DEVELOPMENT STANDARD

QIC-CRF3
Revision C
21 Mar 96

COMMON RECORDING FORMAT FOR USE WITH
FLEXIBLE DISKETTE ENCODED RECORDING FORMATS

For use with the following development standards:

QIC-40-MC
QIC-80-MC
QIC-3010-MC
QIC-3020-MC

Quarter-Inch Cartridge Drive Standards, Inc.

311 East Carrillo Street
Santa Barbara, California 93101
Telephone (805) 963-3853
Fax (805) 962-1541
www.qic.org

(See important notices on the following page)

Important Notices

This document is a development standard adopted by Quarter-Inch Cartridge Drive Standards, Inc. (QIC). This document may be revised several times during the development cycle. It is intended solely as a guide for companies interested in developing products which can be compatible with other products developed using this document. QIC makes no representation or warranty regarding this document, and any company using this document shall do so at its sole risk, including specifically the risks that a product developed will not be compatible with any other product or that any particular performance will not be achieved. QIC shall not be liable for any exemplary, incidental, proximate or consequential damages or expenses arising from the use of this document. This development standard defines only one approach to the product. Other approaches may be available in the industry.

This development standard is an authorized and approved publication of QIC. The underlying information and materials contained herein are the exclusive property of QIC but may be referred to and utilized by the general public for any legitimate purpose, particularly in the design and development of quarter-inch tape cartridge drive subsystems. This development standard may be copied in whole or in part *provided* that no revisions, alterations or changes of any kind are made to the materials contained herein. Only QIC has the right and authority to revise or change the material contained in this development standard, and any revisions by any party other than QIC are totally unauthorized and specifically prohibited.

Compliance with this development standard may require use of one or more features covered by proprietary rights (such as features which are the subject of a patent, patent application, copyright, mask work right or trade secret right). By publication of this development standard, no position is taken by QIC with respect to the validity or infringement of any patent or other proprietary right, whether owned by a Member or Associate of QIC, or otherwise. QIC hereby expressly disclaims any liability for infringement of intellectual property rights of others by virtue of the use of this development standard. QIC has not and does not investigate any notices or allegations of infringement prompted by publication of any QIC development standard, nor does QIC undertake a duty to advise users or potential users of QIC development standards of such notices or allegations. QIC hereby expressly advises all users or potential users of this development standard to investigate and analyze any potential infringement situation, seek the advice of intellectual property counsel, and, if indicated, obtain a license under any applicable intellectual property right or take the necessary steps to avoid infringement of any intellectual property right. QIC expressly disclaims any intent to promote infringement of any intellectual property right by virtue of the evolution, adoption, or publication of any QIC development standard.

Rev. A

- Initial Release from 94-75 Rev F.

Rev. B

- Provides a method of reading and writing the Volume Table.
- Moves format identification to Driver Application Notes.
- Adds a Volume Table Extension entry.
- Adds a Volume Table Unicode Tape Name entry.
- Adds editor's notes for handling unsupported compression and Volume Table extension.

Rev. C

- Provide for CD-ROM Tape format option (Type III).
- Changed multi-partition format to NT style Volume Table Entry (Type II).
- Removed unused filemark options for Type I Volume Table Entry.
- Change driver application notes to allow reading added hidden regions.

TABLE OF CONTENTS

1. SCOPE AND INTRODUCTION

1.1 SCOPE

This Development Standard provides a recording format specification for a cartridge encoded to be compatible with flexible diskette formats. By incorporating the strategies of most streaming tape devices, this Standard creates a sequential storage system independent of the file structures for processing systems, communication systems, and associated equipment. This specification provides a standard code for information interchange independent of these file structures, as agreed upon by the interchange parties.

This Development Standard relates solely to formatting of sequential data as blocks within a single or dual data partition, and allocation of Reed-Solomon Error Correcting Code for data retrieval and interchange. Specific tape drive standards that specify magnetic recording specifications and track geometry may reference this Standard.

1.2 INTRODUCTION

This Development Standard defines the logical format requirements for blocks of data recorded on a 0.25 inch (6.3 mm) and .315 inch (8 mm) wide magnetic tape in a cartridge compatible with QIC-40, QIC-80, QIC-3010 and QIC-3020. The drive recording specification and physical track specifications are defined in these QIC development standards. It is the intent of this Standard to also be compatible with pre-formatted tapes made for QIC-40, QIC-80, QIC-3010 and QIC-3020. The use of this Standard should not preclude use of the QIC-113 logical formats recorded in the native mode for floppy tape. The QIC-113 logical format as defined for streaming tape can also comply with this Standard and thus take advantage of a common low level interface used for SCSI-2, ATAPI and this Standard.

To that end, this Standard dictates automatic generation of the Volume Table entries by the drive or driver. As such, certain modes and settings are defaulted. The multi-cartridge sequence number is always set to one and the compression offset will always start at zero. There is no password support or methods to read or write Data Set names or Source Drive Volume Labels. Although these features can be accommodated by drivers in non-standard ways, their use precludes providing a standard interface. To provide these features with a QIC-113 logical format, the SCSI version of this format should be used. The intent is to provide for backward read capabilities of the floppy native format but to exclude these non-standard features from being used for writing.

This Standard accommodates existing systems using other logical formats for file structures without changing the basic data storage structures. In doing so, this Standard allows much of the hardware dependent code required for data retrieval and storage to be common among these competing systems. This Standard also allows for a common low level interface that is independent of the file structures and hardware. This provides consistent interchange within logical formats. This Standard will recognize two previous styles, QIC-40/QIC-80/QIC-3010/QIC-3020, and NT, and will provide a new Volume Table entry style to accommodate compression techniques used for the QIC CD-ROM compatible format, QCDF. These styles are designated Type I, II, and III respectively. Types II and III provide for dual partitions.

1.3 DEFINITIONS

For the purpose of this Standard, the following definitions apply:

Appendable Point. A point at which a write operation is permitted. This would be following either a BOP, Filemark, or EOD.

Bad Sector. A sector determined to be bad during a read operation.

Block. A group of data bytes recorded by the host application as a unit.

BOP (Beginning of Partition). A partition beginning is defined by Volume Table entires.

BOT (Beginning of Tape) Marker. The BOT Marker is a set of two holes punched in the tape. There are sets of holes provided, the innermost of which is used for identifying the storage position for the cartridge. The additional sets of holes are used to ensure reliability of detection and possibly to encode the tape type.

Note: In the storage position, all of the permissible recording area of the tape is wound on the supply hub and is protected by at least one layer of tape not used for recording data. Cartridges to be interchanged shall be rewound to the storage position prior to interchange.

Byte. A group of 8 data bits operated on as a unit.

Cartridge ID. For cartridges with encoded ID, there are five hole-pair locations between BOT1 and BOT2 that are punched in specific patterns to identify the cartridge. For older cartridges, the distances between marks determine the cartridge type. In either case, the drive or application program interface will identify the cartridge. Reference the QIC-172 Cartridge Type Table.

Compression Extent Header. (Type I) If the Volume Table entry indicates QIC-40/QIC-80 or QIC-3010/QIC-3020 compatible or, for QIC-40/QIC-80, if the QIC-113 Major Signature is Valid and the Compression Used bit is set in the Volume Table entry, then compression will use the Compression Extent Header. For data contained within the Data set, this Quadword header precedes one or more Compression Groups to indicate the next uncompressed byte offset within the data steam from the last Filemark or BOP. The Compression Extent Header never spans segments. If there is less than 18 available bytes remaining within a segment, this region is zero filled to the end of the segment. There is never more that one Compression Extent per segment.

For subsequent linked multi-cartridge backups, this offset is a continuation from a previous Filemark. For the purpose of this standard, multi-cartridge sequence number incrementing and offset continuation is not normally supported but is accommodated to allow backward read compatibility.

Compression Group Header. (Type I) If the Volume Table entry indicates QIC-40/QIC-80 or QIC-3010/QIC-3020 compatible or, for QIC-40/QIC-80, if the QIC-113 Major Signature is Valid and the Compression Used bit is set in the Volume Table entry then compression will use the Compression Group Header. This header is a Word containing the physical size of the data immediately following the Compression Group header. If the Compressed Data Segment Spanning bit is not set in the Volume Table entry then a Compression Group may not span segments for the data contained within the Data Set. The maximum Group size is 63,488 bytes, 62 KB, but the header only indicates the actual number of bytes contained within the current segment. The remainder of the group is determined from the Next Extent Offset in the following segments. When the most significant bit of this header is set, the data within the group is not compressed. If there is less than 18 available bytes remaining within a segment, this region is zero filled to the end of the segment. Otherwise, another compression group is written.

Compression Section Header. (Type III) If the Volume Table entry indicates CRF3, QCDF encapsulation with compression enabled and not QIC-40/QIC-80 or QIC-3010/QIC-3020, then compression will use the Compression Section Header. This is a block aligned 12 byte header that contains a header ID, the physical size of the compressed section, the logical address of the first block within the section, and a word exclusive-or sum followed by the inversion of this value.

CRC (Cyclic Redundancy Code). The CRC is a group of 2 bytes recorded after the end of each sector of data and sector header for the purpose of error detection.

CRF3 Compression Map. (Type III) The Compression Map begins with a CRF3 Compression Map Header followed by CRF3 Compression Map entries. The header and entries contain the descriptions of the location of compressed data and the parameters used when compressing it. Using a separate region makes it easy for software to cache this information in memory and speeds up random access to arbitrary blocks of data within the compressed data partition. The directory partition is always uncompressed.

Data Set. The Data Set is a region of the tape with the beginning and ending segments determined in a Volume Table entry in the Volume Table segment following the Tape Headers. For this

Standard, a Data Set will be a region bounded by either BOP, EOD or Filemarks. If the Directory Last bit is set in the Volume Table entry, then a Filemark is contained within the Data Set at the beginning of the directory.

Deleted Data. The data field within the sector can be written with one of two sync codes. There is a sync code defined as normal and one defined as deleted. A data field written as Deleted must also be read as Deleted to ensure that this information in not accidentally read as normal data.

Density Code. This code identifies the nominal distribution of recorded data information per unit length of track. The density is expressed in kilo-bits per inch (Kbpi). See QIC-172.

ECC (Error Correction Code). Special data used to correct bad data.

ECC Sector. A sector containing ECC data in its data field.

Encoding. A method whereby a group of data bits is translated into a group of recording bits. In this Standard the data is encoded as MFM with the physical format compliant with QIC-40, QIC-80, QIC-3010, and QIC-3020.

EOD (End of Data) Segment. The EOD segment is optionally used to mark the end of the data area. This marker consists of a segment written as Deleted data and may contain a list of any pending defective sectors to be updated in the bad sector list or map. Otherwise, the data is zero filled.

EOT (End of Tape) Marker. The EOT Marker is a single hole punched in the tape to indicate that the usable recording area of the tape has been exceeded, and that the physical end of the tape is approaching. There are three EOT holes to ensure reliable detection.

EW (Early Warning) Marker. The EW Marker is a single hole punched in the tape to indicate the end of the usable recording area in the forward direction.

Filemark. A virtual point, assigned a logical address, within the sequence of data blocks. A read operation will terminate upon reaching a Filemark. Filemarks are always segment aligned for Type I and III and are at the beginning of the Data Set except for the first entry. For Type II, the block at the Filemark address contains undefined data and is not used. For Type II there is a Filemark Segment containing the locations of the Filemarks.

Filemark Segment. (Type II) A segment not contained within a Data Set that begins with a Filemark Segment Header followed by Filemark Segment entries. These entires determine the location of Filemarks.

File Set. Term used in other QIC documents to describe the logical equivalent of the Data Set. A Data Set is recorded as a series of blocks but the File Set is recorded as a stream of bytes of any length. For this standard, the term Data Set will be used.

KBytes (KB). This Standard defines 1 KB to be equal to 1024 bytes.

Logical Early Warning (LEW). An early warning point computed by the drive or application program interface according to a vendor-unique algorithm.

LP (Load Point) Marker. The LP Marker is a single hole punched in the tape to indicate the approaching start of the usable recording area in the forward direction.

Next Extent Offset. (Type I) If the Compressed Data Segment Spanning and Compression Used bit are set in the Volume Table entry then each segment within the Data Set begins with a Word byte offset from the beginning of the segment to the next Compression Extent Header. If there is no Extent header contained within the segment then this offset is zero. If this offset does not equal 2, then the data immediately following the Next Extent Offset is a continuation of the previous Compression Group.

Partition. A sequence of logical data blocks with addresses starting from zero. In this Standard there is either one or two partitions defined.

Reserved. Reserved fields are to be written with zeros and ignored by firmware to facilitate future use by QIC.

Sector. A sector holds 1024 bytes of information identified by a fixed and unique physical ID recorded independent of the data. The physical ID sequence is defined in the QIC-40, QIC-80, QIC-3010, and QIC-3020 specifications.

Segment. The segment is defined as 32 sectors of 1024 bytes. The segments are separated by gaps for drive based navigation. Logical segment numbering starts at zero for the first segment on track zero and are sequential.

Streaming. A method of recording on magnetic tape that maintains continuous tape motion without the requirement to start and stop.

Track. A longitudinal area on the tape along which magnetic signals may be serially recorded.

Underrun. A condition developed when the application transmits data at a rate less than required for streaming.

Vendor Specific. Vendor Specific fields are assigned by QIC for vendors to implement unique features beyond the scope of this document.

1.4 BYTE AND CODE REQUIREMENTS

Byte Length. The data shall be in eight-bit bytes. The 8 bits in each byte are numbered b0 to b7, b7 being the most significant bit.

Byte Order. Unless otherwise stated, all multi-byte value entries are stored least significant to most significant. This is opposite the convention used by SCSI but this is the convention for most systems using these devices. Note: Header information is not normally passed to the application by the drive or driver.

Text Code. Bits b0 to b6 correspond to the 7 least significant bit assignments specified in the American National Standard Code for Information Interchange (ASCII), ANSI X3.4 - 1986. To comply with this Standard, bit 7 shall always be set to Zero and the seven bits b0 through b6 shall represent ASCII characters.

Upon agreement between the interchange parties, other coded character sets may be used. Bit 7 may then be a Zero or a One depending upon the character standards used.

Radix. All numbers are decimal, base 10, unless followed by an 'h' suffix that indicates hexadecimal, base 16.

2. DATA FORMAT

2.1 General Information

Each track is recorded sequentially, starting with track 0, then track 1, and so on. From host application software, data is grouped into blocks of 512 or 1024 bytes. These blocks may then be encapsulated into larger groups for compression. These compression groups may not be aligned on block boundaries unless the CRF3 Signature was detected. For backward read compatibility, if a block is not completed before the logical placement of a Filemark or EOD, then this block is zero padded to completion.

This Standard operates on a segment with sectors and on a block basis. Section 2.2 contains a detailed description of segments, while section 2.3 contains a detailed description of blocks.

Physical Address relates directly to the recorded sector on the tape. Each sector, regardless of its contents, is given a unique physical number. Although a bad sector is not used, it keeps its physical sector number. These sector addresses are used to determine the logical sector number number starting from zero.

The physical address reported by the drive or driver to the application is valid following a Filemark or EOD and will be the logical sector number where reading or writing will proceed. Filemarks and EOD are virtual and occupy no space for Type I and III. For these Types, Filemarks are assigned a logical address and are just before each Data Set except for the first Data Set. Filemarks are also before the segment starting the directory if the Directory Last bit is set in the Volume Table Entry. The segment following EOD will be outside the last Volume Table entry and may be recorded as the EOD segment. For Type II, Filemarks are located using a Filemark Segment and the physical block with the corresponding logical address contains undefined data as a place holder.

Logical Address does not relate to the sectors physically recorded on the tape, but to a numbering system for the logical blocks and Filemarks used by the host. Logical blocks are fixed in length to 1024 bytes for Types I and II. Logical block length for Type III may be 512 or 1024 bytes. Only Data blocks and Filemarks are assigned a sequential logical address from zero starting at the beginning of each partition.

2.2 Segments

General Information. On every track, data is recorded in 1024 byte sectors where 32 sectors constitute a segment that is surrounded with gaps of no data. Therefore, each track consists of sequentially recorded segments and gaps. The segment operation is controlled by the recording drive or driver and is normally invisible to the application. Segments are used to implement Reed-Solomon Error correction and to allow for simplistic positioning.

2.2.1 General Track Layout.

Each Segment contains 3 ECC sectors recorded following the sectors used for data. Sectors previously determined to contain defects are skipped and not used so that at least 4 good sectors must be present for any useful data to be recorded within the segment.

2.2.2 Segment Layout

The number of physically recorded sectors associated with a segment is fixed at 32 but the number of sectors containing data may vary depending on defects detected during certification of the tape and placed within the bad sector map in the Tape Header. The relationship between the data sectors and the ECC sectors as well as the ECC algorithm is defined in the recording specifications.

Table 2-1 General Segment Layout

	bad					bad		
sector 1	sector 2	sector 3	…	sector 28	sector 29	sector 30	sector 31	sector 32
data 0	not used	data 1	data 2-25	data 26	ecc 0	not used	ecc 1	ecc 2

2.2.3 Segment Types

There are 6 different types of segments:

Tape Headers

Volume Table Segment

Data Segments

EOD Segment

Filemark Segment (Type II)

CRF3 Compression Map (Type III)

The first two defect free segments on track 0 contains the Tape Headers followed by the Volume Table segment, a series of Data Segments and optionally followed by the EOD Segment.

Data Segments contain either compressed or uncompressed data blocks. For Type I, Compressed data can optionally span segments and, for multi-cartridge backups, begin with an offset derived from the previous cartridge or cartridges. For that reason, to calculate the number of blocks contained within the first linked and compressed Data Set, the first Compression Extent Header must be read. There is no low level compression for Type II. Type III with QCDF encapsulation uses CRF3 Compression Map and Compression Section Headers to locate and encapsulate compressed block sections.

Segments may be overwritten. To take advantage of this overwrite capability, append operations may begin at EOD (End of Recorded Data) or following any Filemark. To ensure the optional EOD segment is not missed during a read operation, all the sectors are written as Deleted data and may contain a bad sector list of any pending bad sectors to be updated in the Tape Header bad sector map at the time EOD was written. The EOD information is not available to the application and is detected if the majority of the sectors are read as Deleted data. The optional EOD segment will override the EOD indicated within the Volume Table.

2.2.3.1 Tape Header

The Tape Header segment structure is defined in the physical recording specification QIC-40, QIC-80, QIC-3010, or QIC-3020. The Tape Header is repeated and found as the first two segments not written with Deleted data. The Tape Header is automatically updated by the drive or driver.

2.2.3.2 Volume Table

The Volume Table segment shall be created in the format operation following creation of the Tape Header segments as an all “zeros” segment. The Volume Table segment is the first segment in the Logical Area following the Tape Header.

There may be spared segments before the Volume Table to allow for Volume Table Extensions. For Type II Filemark Tables are found before the Data Set. Both Type II and III allow for dual partitions as determined by the Volume Table entry for the Data Set containing the partition.

For Type I, Data Sets that comply to this specification are identified as QIC-113 compatible. The Volume Table will contain an entry for each Data Set stored on the cartridge. More segments for the Volume Table are allocated using the EXVT Volume Table entry. See QIC-80/QIC-3010/QIC-3020.

The first 4 bytes of a used entry shall hold the uppercase ASCII string "VTBL" as a signature. Following all used entries, the end of the Volume Table is indicated by the absence of a signature. For this Standard, only "VTBL" and "EXVT" are used with the "XTBL" and "UTID" entries ignored. If the first entry does not contain a "VTBL" signature, the cartridge is empty. If the last entry slot contains a valid signature other than "EXVT", the cartridge is full. For the purposes of providing backward read capabilities, the QIC-40/QIC-80 and QIC-3010/QIC-3020 Volume Table entries should be handled in the same manner as a Volume Table entry written to this Standard.

Each successive Volume Table entry allocates a segment range beyond the previous with its ending segment number plus one becoming the starting segment number of the next. Any segments skipped before a Data Set entry should be ignored.

Although supported for backward read compatibility, multiple cartridge Volume Table entries are not created. There are considerations in handling multiple cartridge sets for the translation of logical blocks. For these special cases of multiple cartridges, the Volume Table entries for the linked Data Sets are identical except for the Starting and Ending segments, the Multiple Cartridge bit, the Multi-cartridge Sequence Number and Vendor Extension Data.

To be compatible with tapes already containing data from other QIC-compliant Type I applications, available segments are determined by the previous Volume Table entry. A Type III append claims remaining capacity. A tape initialized for Type III can have one, two or three specific Volume Table entries. If using QCDF encapsulation, typically the order would be Compression Map, Partition 1, Partition 0. If the compression map is not used or a single partition implemented, there could be no Compression Map or Partition 1 entry. The use of compression is determined only at initialization of the Type III specific Volume Table entries. If Partition 0 is compressed and the QCDF encapsulation is used, then the Compression Map Volume Table entry is present. Partition 0 will typically start on the first segment of the first even numbered track following Partition 1. The CRF3 Compression Map is not be considered part of the partition.

All character strings shall be left-justified space padded.

Table 2-2 Standard Volume Table Entry (Type I)

Offset	Usage
0-3	Volume Entry Signature. Must equal the ASCII uppercase string “VTBL” for all used entries.

Format Codes 2-5

4-5	Starting Segment Number. Word: Number of first segment within the range allocated to this volume entry (for this cartridge).
6-7	Ending Segment Number. Word: Number of last segment within the range allocated to this volume entry (for this cartridge).

Format Code 6 (QIC-3020 only)

4-7	Number of logical segments used. Represents the total number of logical segments needed to store the data on a volume on a tape. If the volume spans multiple tapes, this number represents only the data on a particular tape.
8-51	Volume Entry Description String. May be any string of ASCII text. (First byte is zero if none specified.) For this Standard, the first byte is zero.
52-55	Volume Entry Storage Data and Time. Doubleword: the bits are encoded as follows: 31-25: Year-1970 (0-127=1970-2097) 24-0: Month, day, hour, minute and second in seconds. Let MO=month (0-11), DY=day (0-30), HR=hour (0-23), MN=minute (0-59), SC=second (0-59), Stored here is SC+60*(MN+60*(HR+24*(DY+31*MO))).
56	Volume Flags. The bits are defined as follows: 0: If zero, then the Volume Table entry is QIC-40/QIC-80 or QIC-3010/QIC-3020 compatible depending on the Tape Header. For QIC-40/QIC-80, if this bit is set and the QIC-113 Major Signature is Valid, then this Volume Table entry is compatible to this and the QIC-113 Standard. This bit is set with QIC-40/QIC-80 and cleared with QIC-3010/QIC-3020 for this Standard. 1: Multiple Cartridge bit. Set if this Data Set spans to another cartridge. This bit is not set for this Standard but must be handled for backward compatible reads. 2: Non-Verification bit. Set if the Data Set was written without verification. 3: Re-direction Inhibit bit. Set if re-direction is disallowed. This bit is not supported by this Standard and is set to zero. 4: Compressed Data Segment Spanning bit. Set if Compression Groups are allowed to span segments. 5: Directory Last bit. This indicates that the File Set Directory section follows the File Set Data Section. For this Standard, this bit indicates the placement of a Filemark before the location of the directory last. This bit is not set if no Filemarks were written. 6-7: Reserved
57	Multi-cartridge Sequence Number. Verifies the correct multiple cartridge loading sequence. The initial cartridge =1, the next = 2, etc. Can be greater than one only in the first Volume Table entry. For this Standard, this number is always set to one.
58-59	QIC-113 Major Signature. Word: Valid if set to 113. For this Standard, set to 113.

60-61 QIC-113 Minor Signature. Word: For this Standard, set to 6.

62-65 CRF3 Signature. For this Standard, the ASCII string “CRF3”

66-67 Minor Version number. 3 for Revision C of this Standard.

68-69 Major Version number. 1 for Compatibility level of this Standard

70-83 Reserved.

84-91 Data Set Password. First byte zero if none specified. For this Standard only non-password Data Sets are read and the password is written as zero.

92-95 Directory Section Size. Doubleword: Number of bytes written past the Filemark. The Filemark will force segment alignment and set the Directory Last bit. If no Filemark is written then this field is zero. As a special case for backward read compatibility, if the Directory Last bit is not set and this field is not zero, then if a block is truncated by this value, it will be zero padded to completion before reading subsequent data.

96-103 Data Section Size. Quadword: Total number of bytes written before the Filemark. For backward read compatibility, on linked Data Sets this includes all cartridges of the multi-cartridge set or if the Directory Last bit is set this includes the sum of the Data Sections up through and including this cartridge. For this Standard, writing a linked Data Set is not supported.

104-105 OS Version Number. The First byte is operating system major version number, second is minor. If undefined, these bytes shall be zero. For this Standard these bytes are zero.

106-121 Source Drive Volume Label. First byte zero if none specified. Otherwise, may be any ASCII string. For this Standard, the first byte is zero.

122 Logical Device File Set Originated From. If undefined, this byte shall be zero. For this Standard this byte is zero.

123 Reserved, set to zero.

124 Byte Compression Method Used. See QIC-123 for definition. The bits are:

- 0-5: Compression code field (3fh indicates vendor specific. Only a code of 01h is allowed for compliance to this Standard.
- 6: Always zero.
- 7: Set to one if compression is used within Data Set.

125 Format and OS Type. Identifies whether the volume data is formatted in Basic DOS or Extended Format. If Extended Format is used, also indicates operating system volume data originated from. The OS type is informational only. The important feature about this field is that if its value is NOT 1, Extended Format is used.

For this Standard, this field is defaulted to the value 1, DOS Basic.

126-127 Reserved, set to zero.

Table 2-3 QIC-40/QIC-80 Volume Table definitions beyond offset 57.

Offset	Usage
58-83	Vendor Extension Data. Reserved for Vendor Unique extensions to the volume entry.
84-91	Data Set Password. First byte zero if none specified. For this Standard only non-password Data Sets are read.
92-95	Directory Section Size. Doubleword: Number of bytes written past the Filemark. The Filemark will force segment alignment and set the Directory Last bit. Thus the exact location during a read operation is determined by moving back to the previous segment boundary just before this range of bytes from the end of the Data Set. If no Filemark is written then this field is zero. As a special case for backward read compatibility, if the Directory Last bit is not set and this field is not zero, then if a block is truncated by this value, it will be zero padded to completion before reading the subsequent data.
96-99	Data Section Size. Doubleword: Total number of bytes written before the Filemark. For backward read compatibility, on linked Data Sets this includes all cartridges of the multi-cartridge set or if the Directory Last bit is set this includes the sum of the Data Sections up through and including this cartridge.
100-101	OS Version Number. The First byte is operating system major version number, second is minor. If undefined, these bytes shall be zero.
102-117	Source Drive Volume Label. First byte zero if none specified. Otherwise, may be any ASCII string.
118	Logical Device File Set Originated From. If undefined, this byte shall be zero.
119	Physical Device File Set Originated From. If undefined, this byte shall be zero.
120	Byte Compression Method Used. See QIC-123 for definition. 0-5: Compression code field (3fh indicates vendor specific. Only a code of 01h is allowed for compliance to this Standard. 6: Always zero. 7: Set to one if compression is used within Data Set.
121-122	OS Type bit mask.
123	ISO Compression Algorithm ID field (values 00h, 01h, or ffh indicates not ISO registered)
124-127	Reserved, set to zero.

Table 2-4 QIC-3010/QIC-3020 Volume Table definitions beyond offset 57.

Offset	Usage
58-83	Vendor Extension Data. Reserved for Vendor Unique extensions to the volume entry.
84-91	Data Set Password. First byte zero if none specified. For this Standard only non-password Data Sets are read.
92-95	Directory Section Size. Doubleword: Number of bytes written past the Filemark. The Filemark will force segment alignment and set the Directory Last bit. Thus the exact location during a read operation is determined by moving back to the previous segment boundary just before this range of bytes from the end of the Data Set. If no Filemark is written then this field is zero. As a special case for backward read compatibility, if the Directory Last bit is not set and this field is not zero, then if a block is truncated by this value, it will be zero padded to completion before reading the subsequent data.
96-103	Data Section Size. Quadword: Total number of bytes written before the Filemark. For backward read compatibility, on linked Data Sets this includes all cartridges of the multi-cartridge set or if the Directory Last bit is set this includes the sum of the Data Sections up through and including this cartridge.
104-105	OS Version Number. The First byte is operating system major version number, second is minor. If undefined, these bytes shall be zero.
106-121	Source Drive Volume Label. First byte zero if none specified. Otherwise, may be any ASCII string.
122	Logical Device File Set Originated From. If undefined, this byte shall be zero.
123	Reserved, set to zero.
124	Byte Compression Method Used. See QIC-123 for definition.
	0-5: Compression code field (3fh indicates vendor specific. Only a code of 01h is allowed for compliance to this Standard.
	6: Always zero.
	7: Set to one if compression is used within Data Set.
125	Format and OS Type. Identifies whether the volume data is formatted in Basic DOS or Extended format. The OS type is informational only. The important feature about this field is that if its value is NOT 1, Extended Format is used.
126-127	Reserved, set to zero.

Table 2-5 Default Volume Description String for Type I

D	A	T	A	_	S	E	T		n	n	n
44h	41h	54h	41h	5fh	53h	45h	54h	20h	30h+	30h+	30h+

The Ns represents a sequential ASCII decimal number with leading zeros starting from 0.

Table 2-6 NT Volume Table entry (Type II)

Offset	Usage
0-3	"VTBL" Volume entry signature
Format Codes 2-5	
4-5	Starting Segment Number of partition for this cartridge
6-7	Ending Segment Number of partition for this cartridge
Format Codes 6 (QIC-3020 only)	
4-7	Number of logical segments used. Represents the total number of logical segments needed to store the data within the partition.
8-51	"Microsoft Windows NT Format 1.0" (left justified blank padded)
52-55	Volume entry storage date and time. Doubleword, the bits are encoded as follows: 31-25: Year-1970 (0-127 = 197- - 2097) 24-0: Month, day, hour, minute and second. Let MO=month (0-11), DY=day (0-30), HR=hour (0-23), MN=minute (0-59), and SC=second (0-59). Stored here is: SC+60*(MN+60*(HR+24*(DY+31*MO)))
56	Volume Flags. The bits are defined as follows: 0: This bit is set for a Type II partition. 1: Multiple Cartridge bit. Set if this Data Set spans to another cartridge. This bit is not set for this Standard. 2: Non-Verification bit. Set if the Data Set was written without verification. 3: Re-direction Inhibit bit. This bit is not supported by this Standard and is set to zero. 4: Compressed Data. Set to zero for this Standard. 5: Directory Last bit.. This bit is not set for Type II partitions. 6-7: Reserved
57	Reserved
58-61	Vendor ID (normally set to "CMS")
62-63	Reserved
64-65	Vendor Software Revision (normally set to 300h)
66-68	Reserved
69	Operating Type and Partition flag Bit 7 set = Directory Partition (1), bits 6-0 = 4 NT
70-91	Reserved
92-95	Segment containing Filemark Segment.
96-99	Partition Size in bytes (locates EOD)
100-127	Reserved

Table 2-7 QCDF Volume Table entry (Type III)

Offset	Length	Usage
0	4	Volume Entry Signature: ASCII uppercase string "VTBL"
Format Code 2-5		
4	2	WORD: Starting Segment Number
6	2	WORD: Ending Segment Number
Format Code 6 (QIC-3020 only)		
4	4	Number of logical segments used. Represents the total number of logical segments needed to store the data within the partition.
8	44	TEXT: Names, left justified space padded uppercase ASCII string "CRF3-COMPRESSION-MAP" or "CRF3-PARTITION-0" or "CRF3-PARTITION-1"
52	4	Volume Entry Storage Data and Time. Doubleword: the bits are encoded as follows: (Ignore if set to zero.) 31-25: Year-1970 (0-127=1970-2097) 24-0: Month, day, hour, minute and second in seconds. Let MO=month (0-11), DY=day (0-30), HR=hour (0-23), MN=minute (0-59), SC=second (0-59), Stored here is SC+60*(MN+60*(HR+24*(DY+31*MO))).
56	1	Volume Flags. The bits are defined as follows: 0: This bit is set for a Type III partition. 1: Multiple Cartridge bit. Set if this Data Set spans to another cartridge. This bit is not set for this Standard. 2: Non-Verification bit. Set if the Data Set was written without verification. 3: Re-direction Inhibit bit. This bit is not supported by this Standard and is set to zero. 4: Compressed Data. Set to zero for this Standard. See offset 66. 5: Directory Last bit.. This bit is not set for Type III partitions. 6-7: Reserved
57	1	BYTE: Multi-cartridge sequence number, set to 0 for this Standard.
58	4	ASCII string “CRF3”
62	4	DWORD: Number of Logical Data Blocks
64	1	Minor Version, starting from 3 for Revision C
65	1	Major Version, starting from 1 For compatibility level of this Standard.
66	1	BYTE: Compression method: 0=uncompressed 1=QIC-122 (see QIC-123) 2-255 reserved
67	1	BYTE: Encapsulation method: 0=no encapsulation 1=standard QIC-113 encapsulaton (segment spanning)

		2=QCDF mapped encapsulation
		3-255=reserved
68	4	BYTE: Partition Flags:
		Bit 0: 0=skip bad segments, 1=skip bad sectors
		Bit 1: 0=512 byte blocks used, 1=1024 byte block used.
		Bit 2-31: reserved, set to zero.
72	56	Unused, Set to zero

Note: As indicated by bit 0 of the Partition Flags not set, all segments with defects are skipped and not used otherwise the normal sector defect mapping is used. Defective sectors and ECC sectors are also excluded in calculations for locating logical blocks. The block addressing for each partition starts at zero and is sequential and contiguous despite exclusions.

Table 2-8 Filemark Segment Header (Type II)

Offset	Length	Usage
0	4	ntm_count. Unsigned number of entries in use. (255 maximum entries.)

Table 2-9 Filemark Segment Entry (Type II)

Offset	Length	Usage
0	4	ntm_type. The type of mark. For this Standard only a value of 1 is used.
4	4	ntm_offset. The unsigned byte offset from the beginning of the partition.

Note: The Filemark Segment entries immediately follow the Filemark Segment Header. Although NT allows for marks other than Filemarks, for the purpose of this Standard, these other marks are not documented. The last entry in the list is identified by a mark with a 0ffffffffh byte offset.

Table 2-10 CRF3 Compression Map Header (Type III)

Offset	Length	Usage
0	4	ASCII String “CRF3”
4	2	WORD: Minor Version: 3 for Revision C of this Standard.
6	2	WORD: Major Version: 1 for compatibility level of this Standard.
8	4	DWORD: Size of Map Header in Bytes (normally 1024)
12	4	DWORD: Valid Bytes in Map
16	4	DWORD: Valid Sections in Map
20	4	DWORD: Logical Blocks Per Section (normally 64)
24	4	DWORD: Bytes Per Group Entry (normally 64)
28	4	DWORD: Offset of First Count in Group Entry (normally 4)
32	4	DWORD: Bytes Per Logical Block (normally 512)
36	988	unused, set to zero

Table 2-11 CRF3 Compression Map Group Entry (Type III)

Offset	Length	Usage
0	4	Physical Block Offset
4	1	Physical Block Count for Section 0
5	1	Physical Block Count for Section 1
		…
	1	Physical Block Count for Section N-1
		(N=Bytes Per Group - Offset of First Count)

Note: CRF3 Compression Map entries immediately follow the CRF3 Compression Map Header.

Table 2-12 Compression Section Header (Type III)

Offset	Length	Usage
0	2	Compressed Segment ID: 7379h
2	2	Count N of Compressed Bytes.
4	4	Logical Block Address of first block within this section.
8	2	16 bit exclusive-or sum of the first 4 words of this header.
10	2	Negated 16 bit exclusive-or sum of the first 4 words of this header.
12	N	Compressed Bytes.

Table 2-13 Extended Volume Table Entry

Offset	Usage

0-3 Volume Entry Signature. Must Equal the ASCII uppercase string “EXVT”

4-5 Parent Segment Number. Word: Number of segment containing this entry.

6-7 Child Segment Number. Word: Number of next segment to extend the parent volume table (for this cartridge).

8-127 Reserved

2.3 Data Blocking

Information in the Volume Table entry determines the type of block being recorded within a Data Set such as compressed, compressed with segment spanning or uncompressed. The valid sectors are determined from the bad sector map in the Tape Header. To support logical blocks from a byte oriented data stream with the inclusion of Filemarks, a translation scheme must be devised. For Type I and II, 1024 byte blocks are always used. For type III, either 512 or 1024 byte blocks may be used and preferably 512 byte blocks. For Type I and III, Filemarks are virtualized just before the segment boundary (segment aligned). For Type II, Filemarks are located by a Filemark Segment preceeding the Data Set and have unused blocks acting as a place holder. Incomplete blocks, as determined by the section size, that precede either a Filemark or EOD will be padded with zero until completion of the block. Data remaining between the end of a section and a Filemark or EOD is ignored. For backward read compatibility, on Data Sets that do not indicate Directory Last and have a non-zero Directory Section Size, the last block in this directory section may require zero padding should the size truncate this block. In this case, once the last block is read and possibly padded, the reading of the Data Section begins at the next byte as this subsequent section is not segment aligned.

For Type I, Filemarks are inserted between Data Sets and at the beginning of the Directory Last section within a Data Set. If no Filemarks are written, there would only be one Data Set without the Directory Last bit set and with the Directory Section Size set to zero.

The optional EOD segment marker is written to mark the End Of recorded Data and to indicate which bad sectors are pending to be updated in the Tape Header at the time this EOD segment was written. The bad sector list, written as Deleted data in the EOD segment, starts at offset 0 with an ASCII signature of “BADSCT”. This signature is followed by a list of 3 byte entries of the logical sector numbers starting with 1. A zero defines the end of the list. If the most significant bit is set for a sector starting a segment, then the entire segment is mapped as bad. This EOD segment provides a safeguard against the lost of data from an incomplete session.

2.3.1 Erase and Append Operations at Appendable Points

Write or erase operations may begin at BOP (the first segment following the Volume Table segment for a write operation) or EOD. Additionally, they may begin at any segment following a Filemark.

Appendable points are created by EOD or a Filemark. Appendable points may be created at the discretion of the application by writing a Filemark.

2.3.2 Data Compression (Type I)

This Standard supports data compression as an optional feature. More than one method of data compression may be employed. This Standard provides for techniques where if the drive is unable to provide for data compression, a low level driver would be able to emulate such capabilities.

The user should also refer to QIC-122 for additional information.

Table 2-14 Example of non-segment spanning Data Set (ECC excluded)

Extent Header (sum of previous bytes from last Filemark or BOP)	Group Header 1 (x)	Data in Group 1 (x bytes)	Group Header 2	Data in Group 2	…	Group Header n	Data in Group n

Table 2-15 Example of first segment in spanning Data Set (segment n ECC excluded)

Next Extent Offset (2)	Extent Header (sum of previous bytes from last Filemark or BOP)	Group Header 1 (x)	Data in Group 1 (x bytes)	Group Header 2	Data in Group 2	Group Header 3 (y)	Data in Group 3 (y bytes)

Table 2-16 Example of segment in spanning Data Set (segment n+1 ECC excluded)

Next Extent Offset (0)	Continuation of Data in Group 3

Table 2-17 Example of segment in spanning Data Set (segment n+2 ECC excluded)

Next Extent Offset (z+2)	Data in end of Group 3 (z bytes)	Extent Header (sum of previous bytes from last Filemark or BOP)	Group Header 4	Data in Group 4	Group Header 5	Data in Group 3

2.3.3 **Data Compression** (Type III - QCDF Encapsulation)

When dealing with compression, the translation used for locating logical blocks is used instead to locate physical blocks that are sequentially addressed relative to the beginning of the data partition. For QCDF encapsulation, an index into the CRF3 Compression Map will provide the physical to logical translation. A number of logical blocks are combined together and compressed to form a Section. The physical starting location and physical sizes for a series of Sections are entered into a

Group Entry. The CRF3 Compression Map Header and CRF3 Compression Group Entries combine together to form the Compression Map.

Data is compressed using the methods described in QIC-122. Using the QCDF encapsulation, the size of a Section is normally 32 KB (64 logical blocks of 512 bytes each) and has a maximum of 62 KB. Sections are zero padded and physical block aligned. If compressed, each Section is block aligned and begins with a Compression Section Header. If there is no savings from compression, then the data is stored uncompressed and the Physical Block Count for this Section is equal to the Logical Blocks Per Section. When decompressed, the data size must equal the size of the Logical Blocks Per Section multiplied by Bytes Per Logical Block.

As Sections are compressed one at a time, the Physical Block Counts are collected in a structure called a Group Entry (normally 64 bytes in size). The first double word of each Group Entry is the Physical Block Offset within the data partition for the first Section. The Physical Block Offsets of other Sections in the Group are obtained by adding the Physical Block Counts of all Sections before Section in question to the Physical Block Offset.

To extract a given logical block, divide the address by the Logical Blocks Per Section. This creates a Section Index and, with the remainder, a logical block offset within that Section. Next, divide the Section index by the number of Sections per Group Entry (normally 60). This creates a Group index and a Group offset. Multiply the Group index by the Group Entry Size and add the Size of the CRF3 Compression Map Header to determine a byte offset into the Map for the desired Group Entry. Adding the Offset of First Count in a Group Entry (normally 4) to the Group offset provides the location of the byte that contains the Physical Block Count. Adding all Section Physical Block Counts to but not including the desired Section to the double word Physical Block offset at the beginning of the Group Entry gives you the physical block address for the Section. The Section can then be read and decompressed, if the count indicates it is compressed, and the data extracted for the logical block indicated by the Section index and logical block offset.

Group Entries beyond the last valid Section may contain garbage data. The Valid Sections in Map field in the Map Header indicates the range of Sections that have been written. Section numbers greater or equal to that value are invalid until they are allocated by writing data to the end of partition. A driver or drive firmware supporting this format would typically buffer an integer multiple of several Group Entries.

3. Driver Application Notes:

1. When appending to a cartridge with the last Data Set not containing a valid CRF3 signature and the write data is not started with a Filemark, start a new Data Set by auto inserting a Filemark at the beginning of this new data. If the previous Data Set was a CRF3 Data Set of the same Type, then start appending within this Data Set. This will ensure non-CRF3 Data Sets are not corrupted.
2. For multiple cartridge linked Data Sets, when directory first is used without compression, every physical data block contained within the first Data Set will be assumed part of this Data Set.
3. When creating a native floppy QIC-113 format Data Set, only directory last can be used. This is done by separating the data section from the directory section with a Filemark. Appending would begin with a Filemark. This would be different from the SCSI version of the QIC-113 format in that the VTS, Volume Table Segment, would be separated from the concatenated data/directory sections with a Filemark. One should consider the native floppy Volume Table a hidden Filemark table and the only VTS in the SCSI version of the QIC-113 format can be used by the application.
4. Read or Write of Volume Table, Tape Headers, Filemark Segment, CRF3 Compression Map:

   This extension does not affect QIC-172. To read or write normally hidden regions of the tape, by using HID (Hidden) bit 7 of byte 5 set to one for either the READ or WRITE commands shall then transfer data in the block size active in the data region of the tape or defaulted to 1 KB blocks. EOD will result after reading or writing the entire space allocated for the selected region.

   The region transferred is selected using the 4 most significant bits of the three byte transfer length with the HID bit set. The region is selected as follows:

Bits 23-20 of Transfer Length	Region
0	Volume Table
1	Tape Header #1
2	Tape Header #2
3	Filemark Segment for Partition 0 (Type II)
4	Filemark Segment for Partition 1 (Type II)
5	CRF3 Compression Map (Type III)
6-fh	Reserved

   The Volume Table region will be automatically extended with an “EXVT” entry. To position within a region, using the LOCATE command with the CP, Change Partition, bit set and partition 255 selected will cause the position information to be held until the hidden operation.

Implementor's Note: *When compressed data is encountered on the medium compliant with QIC that the Drive or Driver cannot decompress, the Drive or Driver should treat the data as uncompressed. In the sense data, the Valid bit, the ILI bit, and the Information field should be set accordingly. (See Data Compression Page in the QIC-157 Standard.)*