

Technical Information Bulletin
for Information Technology -
Sequential Access Partition Management

Secretariat
Information Technology Industry Council (ITI)

Approved January xx, 199x
ASC X3, Information Technology

Abstract

This technical information bulletin provides an interpretation of the use of the sequential access device partition management contained in American National Standard X3.131-1994, Small Computer System Interface-2 (SCSI-2). This additional information is provided to gain a more uniform implementation of the partition management.

Technical Information Bulletin

This interpretation is issued in response to questions which have been raised regarding certain specifications contained in the content of:

ANSI X3.131-1994, Small Computer System Interface - 2

This interpretation was prepared by Technical Committee X3T10, which is responsible for the maintenance of that standard, and was authorized for release by Accredited Standards Committee X3 in order to provide clarifications as quickly as possible in response to questions raised.

This interpretation, while reflecting the technical opinion of the committee responsible for maintaining the standard, is intended solely as supplementary information to users of the standard. The standard is not altered by the issuance of this interpretation. Any subsequent revision, erratum, amendment, or interpretation to the standard may or may not reflect the contents of this interpretation.

CAUTION NOTICE: This interpretation may be revised or withdrawn at any time. The procedures of X3, Information Technology Industry Council require that action be taken periodically to reaffirm, revise, or withdraw this interpretation. Current information is available by calling or writing the Information Technology Industry Council.

CAUTION: The developers of this interpretation have requested that holder's of patents that may be required for the implementation of the interpretation, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this interpretation.

As of the date of publication of this interpretation and following calls for the identification of patents that may be required for the implementation of the interpretation, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any interpretation it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this interpretation.

Published by

Information Technology Industry Council
1250 Eye Street, NW Suite 200
Washington, DC 20005-3922

Telephone (202) 626-5738
Facsimile (202) 638-4922
Email x3sec@itic.nw.dc.us

Copyright ©1995 by Information Technology Industry Council
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of the publisher.

Printed in the United States of America.
ITIC9501/002

Contents

	Page
Foreword	ii
Introduction	iii
1 Scope	1
2 Normative references.....	1
3 Definitions, symbols and abbreviations.....	1
3.1 Definitions.....	1
3.2 Symbols and abbreviations.....	1
4 Request for interpretation.....	1
5 SCSI-2 mode page management review.....	2
5.1 Mode page support.....	2
5.2 Logical unit use of short mode pages.....	2
5.3 Using the tools of SCSI-2 to manage partitions.....	2
6 Medium partition pages(1 - 4).....	2
6.1 Medium partition page(1).....	2
6.1.1 Page length field.....	3
6.1.2 Maximum additional partitions field.....	3
6.1.3 Additional partitions defined field.....	4
6.1.4 FDP, SDP and IDP fields.....	5
6.1.4.1 FDP support.....	6
6.1.4.2 SDP support.....	6
6.1.4.3 IDP support.....	6
6.1.5 PSUM field.....	7
6.2 Medium partition pages(2 - 4).....	7

Tables

1 Medium partition page(1).....	3
2 Short form example of medium partition page(1).....	3
3 Long form example of medium partition page(1).....	4
4 Two partition example of medium partition page(1).....	5
5 Short form two partition example of medium partition page(1).....	5
6 Medium partition pages(2 - 4).....	8

Foreword

This document provides guidance in the use of the sequential access partition page contained in American National Standard X3.131-1994, Small Computer System Interface-2 (SCSI-2). This document does not replace the descriptions in SCSI-2 and is not intended to conflict with SCSI-2. The purpose of this document is to provide more information to gain a more uniform implementation of the SCSI-2 sequential access partition page functions.

This technical information bulletin was developed by Technical Committee X3T10 of Accredited Standards Committee X3 during 1990-1992. The approval process started in 1992 but could not be completed until the publication of X3.131-1994.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be submitted to the X3 Secretariat, Information Technology Industry Council, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This technical information bulletin was processed and approved by Accredited Standards Committee X3, Information Technology. Committee approval of the technical information bulletin does not necessarily imply that all committee members voted for approval.

Introduction

This technical information bulletin provides a non-normative interpretation of use of the sequential access device partition management through mode pages. This technical information bulletin, while reflecting the opinion of the Technical Committee which developed the standard, is intended solely as supplementary information to users of the standard. The standard, ANS X3.131-1994 as approved through the publication and voting procedures of the American National Standards Institute, is not altered by this technical information bulletin. Any subsequent revision to the base standard may or may not reflect the contents of this TIB.

Clause 1 defines the scope of this interpretation

Clause 2 specifies the references

Clause 3 defines the definitions, symbols and abbreviations

Clause 4 contains request for interpretation

Clause 5 contains the general mode page management that relates to this TIB

Clause 6 contains the the four mode pages

Information Technology - Sequential Access Partition Management

1. Scope

A technical information bulletin (TIB) provides an interpretation of a published standard. The interpretation is non-normative and reflects the opinion of the responsible Technical Committee. This TIB in no way alters the mandatory or optional requirements of the base standard. The material contained in this TIB may or may not appear in any subsequent revision to the base standard. See clause 4 for the reason this TIB is published.

2. Normative references

The following standard is the reference for this technical information bulletin.

ANS X3.131-1994, Small Computer System Interface - 2 (SCSI-2).

3. Definitions, symbols and abbreviations

3.1 Definitions

For purposes of this technical information bulletin the following definitions apply.

None.

3.2 Symbols and abbreviations

None.

4. Request for interpretation

The sequential access device command set, clause 10, of the SCSI-2 standard, specifies that support for one partition, partition 0 - the default, is mandatory. Support for more than one partition is optional for any volume (see SCSI-2, 10.1).

This TIB is a response to a request for interpretation on the use of this feature of SCSI-2 and in particular to the proper handling of four mode pages associated with the management of partitions (see SCSI-2, 10.3.3.2 and 10.3.3.3).

Several manufacturers of tape logical units that support partitions have supplied differing implementations in their support of partitions. Some implementations were not compatible with each other and have given rise to various device driver changes to support the different implementations. In particular, the problems have occurred in the implementation of the initiator defined partition (IDP) method. The interpretations of the partition size descriptors, particularly whether or not partition 0 is included, has created a need for this TIB. In addition, there was a question about how mode pages are to be managed in an environment where each logical unit might support a different number of partitions. Also, the means by which an initiator selects the implemented options have been subject of considerable discussion.

This TIB attempts to use the current normative sections of the SCSI-2 standard to provide a basis for a consistent interpretation of the mode pages provided by a logical unit in response to a MODE SENSE command. In addition, this TIB attempts to provide guidance for device driver developers on the use of the same mode pages with the MODE SELECT command.

5. SCSI-2 mode page management review

5.1 Mode page support

SCSI-2 permits management of partitions on sequential access devices through the command set and through mode pages for the MODE SENSE and MODE SELECT commands. SCSI-2 specifies medium partition page(1), page code 11h, as an optional page to manage partitions 0 to 63. Medium partition pages(2 - 4) provided support for partitions 64 to 255 on a single volume. One partition is mandatory and there may be up to 256 partitions on each volume. See SCSI-2, 10.3.3.2 and 10.3.3.3.

These mode pages are managed using the MODE SENSE and MODE SELECT commands and are subject to the rules for providing pages to an initiator and interpreting pages sent from an initiator. If any of these mode pages are supported, they are provided in the default form, the changeable form, and the current form. Support for the saved page form is optional. See SCSI-2, 8.2.8, 8.2.10.1 through 8.2.10.4.

If more than one partition is supported, medium partition page(1) is supported to report the level and type of support by a logical unit.

A mode page sent to a logical unit in a MODE SELECT command is required to be the same length as the same page returned in a MODE SENSE command response. By inference, a mode page returned by a MODE SENSE command is required to be constant, for the logical unit. Otherwise, the initiator cannot reliably determine the correct length of the page to send at any given time. See SCSI-2, 8.2.8, 8.2.10 and 8.3.3.

5.2 Logical unit use of short mode pages

A logical unit may have short page lengths if the mode page ends on byte boundaries, and more specifically, on field boundaries before the end of the page specified in the SCSI-2 standard. The short form of a mode page usually means that any corresponding features identified in the missing fields of the mode page are not implemented, or that the value of a field is considered to be zero, or that the field is not used by the logical unit. In any case, the missing parameters are non-changeable. See SCSI-2, 8.3.3, in the mode page format description.

5.3 Using the tools of SCSI-2 to manage partitions

The SCSI-2 rules for mode page management give a logical unit several compliant mechanisms for supporting mode pages and thereby imply several rules for how initiators interpret mode pages. Operationally, the mode page management rules (see SCSI-2, 8.2.8 and 8.3.3) specify that an initiator may not change the page length from that reported in the MODE SENSE command response for the same mode page. If the initiator does attempt to change the length, it should expect to receive a CHECK CONDITION status for any attempt to send a page with a different page length (SCSI-2, 8.3.3). Attempts to set values in parameters outside implemented ranges also result in CHECK CONDITION status either to reject the command for having invalid parameter data or, for values within a valid range, to report that a rounding operation has been used by the logical unit on one or more parameters in a mode page. See SCSI-2, 6.5.4 and 8.2.8.

6. Medium partition pages(1-4)

6.1 Medium partition page(1)

Table 1 shows the parts of medium partition page(1) that relate to this TIB. See SCSI-2, 10.3.3.2. Fields of no interest to this TIB are marked with "XX...XX."

Table 1 - Medium partition page(1)

Bit Byte	7	6	5	4	3	2	1	0	
0	xx	xxxxxx	Page code (11h)						
1	Page length								
2	Maximum additional partitions								
3	Additional partitions defined								
4	FDP	SDP	IDP	PSUM		xxxxxx			
5	xxxxxx								
6	xxxxxx								
7	xxxxxx								
Partition size descriptors									
0	(MSB)			Partition					
1				Size				(LSB)	

6.1.1 Page length field

The page length field identifies the number of bytes that follow in each mode page and can help determine the level of support in a mode page when a mode page is supported. The examples in clauses below show how the page length field may be used. See SCSI-2, 8.2.8 and 8.3.3.

6.1.2 Maximum additional partitions field

The maximum additional partitions field indicates the maximum number of additional partitions supported by the logical unit and that could exist on a volume. A value of zero indicates that no additional partitions are present or allowed. A default partition, partition(0) is required, so the minimum number of partitions supported is one per volume. This field name contains the word "additional" since all volumes have the required default partition(0). When the value of this field is greater than zero, it indicates that multiple partitions per volume are supported. It does not mean that they are present on the currently mounted volume (see 6.1.3).

Since there is no provision in a partition size descriptor to identify the number of a partition, the name of a partition is the relative position of the partition size descriptor in medium partition pages(1-4). The size of partition(0), if specified, is specified in the first partition size descriptor of medium partition page(1), that is in bytes(9-10). Partition 1 is specified in bytes(11-12), etc. The other mode pages are used when more than 64 partitions are supported on a volume. If additional partitions are supported and information about the partitions is to be provided, there is one partition size descriptor following the partition size descriptor for partition 0 for each valid partition.

Tables 2 and 3 show two variations for reporting a single partition per volume. Table 2 shows an example of the short form of medium partition page(1) where the logical unit declares a fixed number of partitions (just the default) and has chosen not to declare the estimated size of the default partition. Note that PSUM may be any valid value since there is no field that uses it.

Table 2 - Short form example of medium partition page(1)

Bit Byte	7	6	5	4	3	2	1	0
0	xx	xxxxxx	Page code (11h)					
1	Page length = 06h							
2	Maximum additional partitions = 00h							
3	Additional partitions defined = 00h							
4	FDP=1	SDP=0	IDP=0	PSUM=xx		xxxxxx		
5	xxxxxx							
6	xxxxxx							
7	xxxxxx							

This format may be used when the sequential access density code in the density code field of the mode parameter header is considered sufficient to describe the estimated size of the default partition of a volume. It is recommended that the estimated size be present, but it is not required (see table3).

Table 3 - Long form example of medium partition page(1)

Bit Byte	7	6	5	4	3	2	1	0
0	xx	xxxxxx	Page code (11h)					
1	Page length = 08h							
2	Maximum additional partitions = 00h							
3	Additional partitions defined = 00h							
4	FDP=1	SDP=0	IDP=0	PSUM=10b		xxxxxx		
5	xxxxxx							
6	xxxxxx							
7	xxxxxx							
8	(MSB)		Partition size = 07D0h					
9	(2000 decimal)						(LSB)	

Table 3 shows an example of the long form of medium partition page(1) where the logical unit declares a fixed number of partitions and includes an estimate of the size of that partition. This is the recommended form of this mode page.

In each case, these mode page fields would be marked non-changeable and the current and default mode page content would be identical.

There is no guarantee that the size estimated from the density code field or from a partition size field can be achieved; it is an estimate only. This value normally represents a nominal capacity, excluding compression.

For a multiple fixed partition example, see 6.1.3.

6.1.3 Additional partitions defined field

When only the default partition is supported, both the maximum additional partitions field and the additional partitions defined field are set to zero. See examples in 6.1.2. When additional partitions are allowed (maximum additional partitions > 0), the value of this field plus one identifies the number of partitions that actually exist on a mounted volume (i.e., the number of partition size descriptors, if present, that have values greater than zero in any valid unit of measure). This field has no meaning when the logical unit is not ready.

Table 4 shows an example of the medium partition page(1) for a volume that has two fixed defined partitions (FDP = 1b) and therefore two logical unit-determined, fixed sized partitions (not necessarily the same size). These are considered non-changeable fields and the current and default mode page contents would be identical.

Table 4 - Two partition example of medium partition page(1)

Bit Byte	7	6	5	4	3	2	1	0
0	xx	xxxxxx	Page code (11h)					
1	Page length = 0Ah							
2	Maximum additional partitions = 01h							
3	Additional partitions defined = 01h							
4	FDP=1	SDP=0	IDP=0	PSUM=10b		xxxxxx		
5	xxxxxx							
6	xxxxxx							
7	xxxxxx							
8	(MSB)	Partition size = 03E8h						(LSB)
9	(1000 decimal)							
10	(MSB)	Partition size = 03E8h						(LSB)
11	(1000 decimal)							

Since there is one additional partition defined, two partition size descriptors are present in the example of this mode page. Bytes(8-9) contain the partition size descriptor for the default partition(0). Bytes(10-11) contain the partition size descriptor for the additional partition(1). In this example, the two partitions are the same size, but there is no requirement that they be the same size. The page length is 0Ah.

Just as in the single partition example in 6.1.2, it is not mandatory to declare the partition sizes. However, reporting partition sizes using the partition size descriptors is recommended. If the short form were to be used, the page length would be 06h. In this case, the initiator could not determine the size of either partition, but it could determine that two partitions exist. See table5 for an example of this short form.

Table 5 - Short form two partition example of medium partition page(1)

Bit Byte	7	6	5	4	3	2	1	0
0	x	xxxxxx	Page code (11h)					
1	Page length = 06h							
2	Maximum additional partitions = 01h							
3	Additional partitions defined = 01h							
4	FDP=1	SDP=0	IDP=0	PSUM=xx		xxxxxx		
5	xxxxxx							
6	xxxxxx							
7	xxxxxx							

6.1.4 FDP, SDP and IDP fields

A fixed data partitions (FDP) field of 1b indicates that the logical unit assigns partitions based on its internal fixed definition of partitions. A select data partitions (SDP) field of 1b indicates that the logical unit divides the medium into the number of partitions plus 1 as specified by the additional partitions defined field when supplied by an initiator in the mode page in a MODE SELECT command. An initiator-defined partitions (IDP) field of 1b indicates that the initiator defines the number of additional partitions and size of each partition using the additional partitions defined field and the partition size descriptors.

Each logical unit must declare one method for each mounted volume. It may not be possible for a logical unit to declare these attributes until a volume is mounted and the logical unit is ready. The characteristic that determines the level of partition support may be contained on the volume and may not be an intrinsic attribute of the logical unit. This is similar to other removable medium devices that can support read only, write once, and rewriteable medium. The determining factor is the volume mounted and not the device itself. Some devices may only be capable of one type and level of support and can declare its capability at any time. However, it is best to depend on the parameters returned in MODE SENSE when the logical unit is ready to be the most accurate reflection of support at that time.

6.1.4.1 FDP support

The FDP field, when 1b, indicates that the logical unit manages the number, placement, and size of partitions on a volume. If the maximum additional partitions field is greater than zero, then that number of additional partitions, beyond the default partition, exist. Their size and location are controlled by the logical unit. The additional partitions defined field has a value equal to the maximum additional partitions field value to indicate that they all exist. Some examples of medium partition page(1) have been given in previous tables.

The PSUM field has a valid value. The partition size descriptors, if present, specify the estimated size of each partition. When FDP = 1b, each partition size descriptor has a value greater than zero. SDP and IDP are each 0b. It is recommended that the partition size descriptors be present to give an estimate of the size of each partition since they are not required to be the same size.

Since the size of the medium partition page(1) is required to be constant for a given logical unit (see 5.1), support of IDP = 1b and FDP = 1b requires that the partition size descriptors be present when FDP = 1b. For the MODE SELECT command, the partition size descriptors are ignored. For the MODE SENSE command, the partition size descriptors specify the estimated size of each partition.

6.1.4.2 SDP support

The SDP field, when 1b, indicates that the logical unit manages the partitions on a volume and therefore the size and placement of each partition based upon a count supplied by an initiator. The initiator controls the number of partitions that exist (greater than or equal to one). The maximum additional partitions field (n, where n is greater than zero) indicates that partition 0 plus up to n additional partitions may be defined on a volume. An initiator may specify any number of partitions to be placed on the volume from 1 (the default and minimum) to a maximum of n + 1 partitions by using the additional partitions defined field.

In a MODE SELECT command, an initiator sets the additional partitions defined field in medium partition page(1) to the number of additional partitions desired (m, where m is less than or equal to n). The number of currently valid partitions then becomes m+1 upon successful completion of the MODE SELECT command.

In MODE SENSE commands, the additional partitions defined field has a value equal to the actual number of additional partitions defined on the volume (m). The value in this field is not meaningful when the logical unit is not ready. The PSUM field has a valid value. The partition size descriptors, if present, specify the estimated size of each partition (i.e., greater than zero). FDP and IDP are each 0b.

It is recommended that the partition size descriptors be present to give an estimate of the size of each partition since they are not required to be the same size and the size of each partition is selected by the logical unit. Since all of these partitions are defined, each partition size descriptor, when present, contains a value greater than zero.

As with the case of FDP = 1b (see 'FDP support'), support of both IDP = 1b and SDP = 1b requires that the partition size descriptors be present when SDP = 1b. For the MODE SELECT command, the partition size descriptors are ignored. For the MODE SENSE command, the partition size descriptors specify the estimated size of each partition.

When m is less than n it is likely, but not required, that the partitions would be assigned partition numbers from 0 to m. If the partition numbers are not assigned in this manner, then the logical unit should provide Medium Partition mode page support for (m+1) partition size descriptors with a value greater than zero to identify which are the valid partitions at the present time. If this type of support is present, the logical unit should support (n+1) partition size descriptors and provide for the entire set of possible partitions.

6.1.4.3 IDP Support

The IDP field, when 1b, indicates that an initiator decides the number of and the size of each partition on a volume; the logical unit manages the placement of partitions on the volume. There is no requirement that the partitions be placed in a specific order on a volume. It is sufficient for the logical unit to be able to determine that a partition exists and to position to the partition if requested. With IDP = 1b support it is possible that some of the space on a volume is not allocated to existing partitions. This may permit a logical unit to add

additional defined partitions at a later time without disturbing currently existing partitions. Such support in a logical unit is optional.

The minimum number of partitions is 1, the default partition. The maximum additional partitions field (n, where n is greater than or equal to zero) indicates that partition 0 plus up to n additional partitions may be defined on the volume by an initiator.

An initiator may specify any number of partitions to be placed on the volume from 1 (the default) to a maximum of (n+1) partitions by using the additional partitions defined field. In addition, the initiator specifies the size of each partition using the partition size descriptors.

This means that the logical unit provides mode page support for (n+1) partition size descriptors in the medium partition pages(1-4). Partitions that are currently undefined have their corresponding partition size descriptor value set to zero.

The initiator chooses the number of additional partitions in the additional partitions defined field (m, where m is less than or equal to n) in the medium partition page(1). The estimated size of each partition is specified in the partition size descriptors, including the default partition. Since the standard requires a default partition to be present on each volume, the partition size descriptor value for partition zero is greater than zero, although it can be of minimal size based on the PSUM field value.

In MODE SENSE commands, the additional partitions defined field has a value equal to the number of additional partitions currently on the volume (m) and the value is less than or equal to the maximum additional partitions field value (n). The PSUM field has a valid value.

There are n+1 partition size descriptors in both the MODE SENSE mode pages and the same number of partition size descriptors in the MODE SELECT mode pages. There are m + 1 partition size descriptors with a value greater than zero. FDP and SDP are each 0b.

Partition zero is required to exist with a non-zero size on each volume and therefore has a partition size descriptor value greater than zero. Since the maximum number of partition size descriptors is present (n+1) in the pages, some of the descriptors, except the first one in mode page 11h, will have a zero value when m < n. When m = n, all partition size descriptor values have a value greater than zero.

6.1.5 PSUM field

The partition size unit of measure (PSUM) field specifies three units in which the partition size descriptors identify or specify a partition size. The values 00b, 01b, and 10b define the units as bytes, kilobytes, and megabytes, respectively. The value 11b is reserved.

The bytes unit of measure is 10^{**0} (or one) and the megabytes measure is 10^{**6} (or 1 000 000). The measure for megabytes is derived from the same unit of measure as the megatransfers per second on the physical bus. There, units are measured using the minimum period for the synchronous data transfer period which is 100 ns or 10 Mhz or 10^{**6} hz. With these two units of measure for PSUM having been defined, that leave only kilobytes to be defined. That unit of measure is 10^{**3} (1 000) for symmetry with the other two defined units of measure.

6.2 Medium partition pages (2 - 4)

The sequential access device mode pages with page codes 12h, 13h, and 14h are defined as medium partition pages 2, 3, and 4, respectively. See SCSI-2, 10.3.3.3. Each page contains up to 64 partition size descriptors. See table 6.

Table 6 - Medium partition pages (2 - 4)

Bit Byte	7	6	5	4	3	2	1	0	
0	xx	xxxxxx	Page Code (12h or 13h or 14h)						
1	Page length								
Partition size descriptors									
0	(MSB)	Partition							
1	Size						(LSB)		

If one of these medium partition pages is supported, each lower numbered medium partition page is supported and each at maximum length. That is, if medium partition page(3) is supported then at least 129 partitions are supported on a volume (64 for medium partition page(1), 64 for medium partition page(2), and at least one from medium partition page (3)).

If a sequential access logical unit supports only FDP = 1b, these pages are not required to be supported if the logical unit does not intend to disclose the estimated size of the partitions. Such an implementation is not recommended, but it is permissible under the rules for short mode page support.

In the case of SDP = 1b, the logical unit may not support these mode pages, but such support is recommended.

When IDP = 1b is supported and the maximum additional partitions field value in Medium Partition page(1) is greater than 63, one or more of these mode pages is required to be supported so that there is one partition size descriptor for each partition supported(1 + maximum additional partitions).