

To: T10 Technical Committee  
 From: Rob Elliott, HP (elliott@hp.com)  
 Date: 5 March 2004  
 Subject: 04-082r0 SBC-2 Replace Notch and Partition mode page with READ CAPACITY

### Revision history

Revision 0 (5 March 2004) First revision (split off from 04-075r0)

### Related documents

sbc2r12 - SCSI Block Commands - 2 revision 12

### Overview

The **Notch and Partition mode page** contains these fields:

- a) NOTCHED DEVICE bit - 1 means the medium is notched
- b) LOGICAL OR PHYSICAL NOTCH bit - 0 means boundaries use cylinder and head, 1 means they use LBAs
- c) MAXIMUM NUMBER OF NOTCHES field - 2 bytes (unchangeable)
- d) ACTIVE NOTCH field - 2 bytes - defines the notch to which all mode page accesses apply
- e) STARTING BOUNDARY field - 4 bytes (unchangeable) - 4 byte LBA if LPB=1, 3 byte cylinder and 1 byte head if LPB=0
- f) ENDING BOUNDARY field - 4 bytes (unchangeable) - 4 byte LBA if LPB=1, 3 byte cylinder and 1 byte head if LPB=0
- g) PAGES NOTCHED field - 8 bytes (unchangeable) - bit map of mode pages that may be different for different notches

Issues:

- a) Cylinder/head/sector values are rapidly becoming obsolete in the rest of SBC-2. If LOGICAL OR PHYSICAL NOTCH bit is set to 0, the STARTING BOUNDARY field and ENDING BOUNDARY field are not big enough to contain real values for a >2 TB medium and would contain fake values.
- b) The STARTING BOUNDARY field and ENDING BOUNDARY field do not support 8 byte LBAs
- c) The PAGES NOTCHED field introduces the possibility that mode pages don't apply to the whole logical unit, just the "active notch," a mode page policy not comprehended by SPC-3. It doesn't support mode subpages. The usefulness for this field seems to be gone. The Format Device mode page, particularly the SECTORS PER TRACK field, used to be the key mode page affected by ACTIVE NOTCH field, but that mode page is obsolete in SBC-2. Some parallel SCSI drives also marked the Disconnect-Reconnect mode pages so the READ BUFFER FULL RATIO and WRITE BUFFER EMPTY RATIO fields could be zone specific; those fields are not very useful any more.
- d) The industry calls these zones, not notches

Only these fields make sense:

- a) ZONED DEVICE bit - the medium is zoned or not
- b) MAXIMUM NUMBER OF NOTCHES field - 2 bytes (unchangeable)
- c) STARTING BOUNDARY field - 8 bytes - 8 byte LBA
- d) ENDING BOUNDARY field - 8 bytes - 8 byte LBA

Rather than a mode page requiring this algorithm to access:

```

for i=0 to Maximum Number of Notches {
    set the Active Notch=i
    read the Starting Boundary (and Ending) Boundary
}

```

the drive could just return a fixed data structure listing all the zone boundaries. As read-only, medium dependent information, this is best retrieved with READ CAPACITY (16) command using a new field in the CDB to request zone information rather than capacity information. It would return:

- a) ZONED MEDIUM bit (1 bit in a 2 byte field)
- b) ZONE BOUNDARY LIST LENGTH field (2 bytes)
- c) zone boundary 0 (8 bytes) (not LBA zero; the first ending boundary)

- d) zone boundary 1 (8 bytes)
- e) zone boundary 2 (8 bytes)
- f) ...
- g) zone boundary n (8 bytes)(the capacity)

### Suggested changes

**3.1.28 notch:** All or part of the medium having a consistent set of geometry parameters. Notches are used to increase storage capacity by optimizing the number of bytes per track between the inner and outer tracks.

**3.1.28 zone:** All or part of the medium having a consistent set of performance parameters. Also known as a notch.

### 4.11 Notched devices

A notched (also known as zoned) device has areas of the medium with geometry changes. In the simplest case, the entire medium consists of a single notch. Multiple notches are often used to increase capacity of the device. On a disk, the inner tracks are physically shorter than the outer tracks. As a result, if each track is made to store the same number of data bits, the data is packed more densely on the inner tracks than the outer tracks. By using notches, the outer tracks may contain a different number of sectors than the inner tracks, while balancing the data density. This results in increased capacity.

### 5.16 READ CAPACITY (16) command

The READ CAPACITY (16) command (see table 1) provides a means for the application client to request information regarding the capacity of the block device. This command is implemented as a service action of the SERVICE ACTION IN opcode. This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.7)

**Table 1 — READ CAPACITY (16) command**

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (9Eh)							
1	<a href="#">MEDIUM INFORMATION TYPE</a>			SERVICE ACTION (10h)				
2	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
9								
10	(MSB)	ALLOCATION LENGTH						(LSB)
13								
14	Reserved							PMI
15	CONTROL							

The [MEDIUM INFORMATION TYPE](#) field specifies the format of the parameter data returned (see table 2).

**Table 2 — Medium information type**

<a href="#">Value</a>	<a href="#">Description</a>
<a href="#">000b</a>	<a href="#">Long read capacity data (see table 4)</a>
<a href="#">001b</a>	<a href="#">Zone data (see table 3)</a>
<a href="#">All others</a>	<a href="#">Reserved</a>

See the LOCK UNLOCK CACHE (10) command (see 5.4) for a definition of the LOGICAL BLOCK ADDRESS field. See the READ CAPACITY (10) command (see 5.12) for a description of the other fields in this command.

The zone data is defined in table 3.

**Table 3 — Zone data**

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							ZONED MEDIUM
1	Reserved							
2	(MSB)	ZONE BOUNDARY LIST LENGTH (n - 3)						(LSB)
3								
4	(MSB)	ZONE ENDING BOUNDARY LOGICAL BLOCK ADDRESS (first)						(LSB)
11								
...	...							
n - 7	(MSB)	ZONE ENDING BOUNDARY LOGICAL BLOCK ADDRESS (last)						(LSB)
n								

A ZONED MEDIUM bit set to one indicates that the medium is zoned. A ZONED MEDIUM bit set to zero indicates that the medium is not zoned.

The ZONE BOUNDARY LIST LENGTH field indicates the length of the zone boundary list that follows.

Each ZONE ENDING BOUNDARY LOGICAL BLOCK ADDRESS field indicates the LBA of a logical block that ends a zone boundary. The zone boundary list shall be sorted in ascending order. The first zone always starts at LBA zero, and the first list entry is the ending LBA of that zone. The last zone always ends at the capacity of the medium and is indicated by the last list entry. If the medium has only one zone, only one list entry is present.

The long read capacity data is defined in table 4.

**Table 4 — Long read capacity data**

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)	RETURNED LOGICAL BLOCK ADDRESS						(LSB)
7								
8	(MSB)	BLOCK LENGTH IN BYTES						(LSB)
11								
12	Reserved							PROT_EN
13								
31	Reserved							

The RETURNED LOGICAL BLOCK ADDRESS field and BLOCK LENGTH IN BYTES field of the long read capacity data are the same as the in the short read capacity data described in the READ CAPACITY (10) command (see 5.12). The maximum value that shall be returned in the RETURNED LOGICAL BLOCK ADDRESS field is FFFFFFFF FFFFFFFEh.

A PROT\_EN bit set to one indicates that the medium was formatted with protection information (see 4.15) enabled. A PROT\_EN bit set to zero indicates that the medium was not formatted with protection information enabled.

### 6.3.1 Mode parameters overview

...

The mode page codes for direct-access devices are shown in table 5.

**Table 5 — Mode page codes for direct-access devices**

Mode page code	Description	Reference
00h	Vendor-specific (does not require page format)	
01h	Read-Write Error Recovery mode page	6.3.4
02h	Disconnect-Reconnect mode page	SPC-3
03h	Obsolete (Format Device mode page)	
04h	Obsolete (Rigid Disk Geometry mode page)	
05h	Obsolete (Flexible Disk mode page)	
06h	Reserved	
07h	Verify Error Recovery mode page	6.3.5
08h	Caching mode page	6.3.2
09h	Obsolete	
0Ah	Control mode page	SPC-3
0Bh	Obsolete (Medium Types Supported mode page)	
0Ch	<a href="#">Obsolete</a> (Notch And Partition mode page)	<del>6.3.3</del>
0Dh	Obsolete	
0Eh - 0Fh	Reserved	
10h	XOR Control mode page	6.3.6
11h - 17h	Reserved	
18h	Protocol-Specific LUN mode page	SPC-3
19h	Protocol-Specific Port mode page	SPC-3
1Ah	Power Condition mode page	SPC-3
1Bh	Reserved	
1Ch	Informational Exceptions Control mode page	SPC-3
1Dh - 1Fh	Reserved	
20h - 3Eh	Vendor-specific (does not require page format)	
3Fh	Return all mode pages (valid only for the MODE SENSE command)	SPC-3

In some cases the mode pages do not apply to the entire logical unit. This is controlled by the Notch And Partition mode mode (see 6.3.3).

### 6.3.3 Notch And Partition mode page

The Notch And Partition mode page (see table 91) contains parameters for direct access devices that implement a variable number of blocks per cylinder and support this mode page. Each section of the block device with a different number of blocks per cylinder, than other sections, is referred to as a notch.

The parameters savable (PS) bit is only used with the MODE SENSE command. This bit is reserved with the MODE SELECT command. A PS bit set to one indicates that the device server is capable of saving the modepage in a non-volatile vendor-specific location.

A notched device (ND) bit set to zero indicates that the block device is not notched and that all other parameters in this mode page shall be returned as zero by the device server. A ND bit set to one indicates that the block device is notched. For each supported active notch value this mode page defines the starting and ending boundaries of the notch.

A logical or physical notch (LPN) bit set to zero specifies that the boundaries are based on the physical parameters of the block device. The cylinder is considered most significant, the head least significant. A LPN bit set to one specifies that the notch boundaries are based on logical blocks of the block device.

The MAXIMUM NUMBER OF NOTCHES field indicates the maximum number of notches supported by the logical unit. This field shall be reported as unchangeable.

The ACTIVE NOTCH field indicates the notch to which this and subsequent MODE SELECT and MODE SENSE commands shall refer, until the active notch is changed by a subsequent MODE SELECT command. The value of the active notch shall be greater than or equal to 0000h and less than or equal to the maximum number of notches. An active notch value of zero indicates that this and subsequent MODE SELECT and MODE SENSE commands refer to the parameters that apply across all notches.

The STARTING BOUNDARY field indicates the beginning of the active notch or, if the active notch is zero, the beginning boundary of the logical unit. If the LPN bit is set to one, then the four bytes represent an LBA. If the LPN bit is set to zero, then the three most significant bytes shall represent the cylinder number and the least significant byte shall represent the head number. This field shall be reported as unchangeable. When used with the MODE SELECT command this field is ignored.

The ENDING BOUNDARY field indicates the ending of the active notch or, if the active notch is zero, the ending of the logical unit. If the LPN bit is set to one, then the four bytes represent an LBA. If the LPN bit is set to zero, then the three most significant bytes shall represent the cylinder number and the least significant byte shall represent the head number. This field shall be reported as unchangeable. When used with the MODE SELECT command this field is ignored.

Each notch shall span a set of consecutive logical blocks on the block device, the notches shall not overlap, and no logical block shall be excluded from a notch.

The PAGES NOTCHED field is a bit map of the mode page codes that indicates mode pages that contain parameters that may be different for different notches. The most significant bit of this field corresponds to PAGE CODE 3Fh and the least significant bit corresponds to PAGE CODE 00h. If a bit is set to one, then the corresponding mode page contains parameters that may be different for different notches. If a bit is set to zero, then the corresponding mode page contains parameters that are constant for all notches. This field shall be reported as unchangeable.