

CONGRUENT SOFTWARE, INC.
98 Colorado Avenue
Berkeley, CA 94707
(510) 527-3926
(510) 527-3856 FAX

FROM: Peter Johansson
TO: T10 SBP-3 working group
DATE: June 26, 2001
RE: RBC-2 commands for extent management

This document contains the core of a set of commands that could be used by RBC-2 devices to manage and access "extents", regions of contiguously addressable logical blocks within a device's medium.

Extents are, in effect, dynamic partitions managed by the device instead of by data structures visible to the initiator. Some of the extents may be identified as accessible *via* a protocol foreign to RBC, e.g., AV/C disks. Methods exist to make extents accessible to RBC while simultaneously preventing access by the foreign protocol and to subsequently return all control over the extent to the other protocol.

Where existing text in RBC would be affected, such as in a table or in a command, the modifications are shown with change bars. For entirely new commands, only the material changed from one revision of this document to the next is highlighted with change bars.

The 01-180r1 revision incorporates the following changes:

- a) The BLOCK ADDRESS field in the read, write and verify commands was increased to 48 bits;
- b) The EXTENT ID field was increased to 32 bits;
- c) A "window" capability was added to the EXTENT DIRECTORY command to permit piecewise retrieval of very large quantities of extent information;
- d) Extents, when created, are initially locked for exclusive RBC access and must be explicitly unlocked before they are available to other protocols supported by the device; and
- e) The READ CAPACITY command was modified to permit the amount of unallocated space to be queried.

5 Reduced Block Commands

The Reduced Block Command set (RBC-2) for block device logical units is shown in Table 2.

Table 2 - Reduced Block Command set

Command name	OpCode	Command Support		Reference
		Fixed	Removable	
EXTENT DIRECTORY	XXh	O	O	RBC-2
EXTENT MANAGEMENT	XXh	O	O	RBC-2
QUERY EXTENT	XXh	O	O	RBC-2
READ EXTENT-RELATIVE	XXh	O	O	RBC-2
VERIFY EXTENT-RELATIVE	XXh	O	O	RBC-2
WRITE EXTENT-RELATIVE	XXh	O	O	RBC-2
Command Support key: M = support is mandatory; O = support is optional.				

The CONTROL byte (the last byte of the CDB) shall be set to zero.

5.0a EXTENT DIRECTORY command

The EXTENT DIRECTORY command (see Table 2a) provides a means for the initiator to request the directory of extents managed by the RBC device. An extent is a logically contiguous range of blocks identified by an EXTENT ID.

Table 2a - EXTENT DIRECTORY Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (XXh)							
1	(MSB)							
2								
3	EXTENT DIRECTORY OFFSET							
4								(LSB)
5	(MSB)							
6								
7	ALLOCATION LENGTH							
8								(LSB)
9	CONTROL = 00h							

The EXTENT DIRECTORY OFFSET and ALLOCATION LENGTH fields permit the return of less than the maximum 536,870,912 bytes of extent information that may be available. EXTENT DIRECTORY

data may be returned starting from the byte specified by EXTENT DIRECTORY OFFSET for a total of ALLOCATION LENGTH bytes.

EXTENT DIRECTORY data (see Table 2b) shall be returned to the initiator prior to the return of GOOD status for the command.

Table 2b- EXTENT DIRECTORY data

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
	DEFAULT RBC EXTENT							
1								(LSB)
2	(MSB) _____							
	DATA LENGTH							
3								(LSB)
4	(MSB) _____							
	EXTENT DIRECTORY							
n								(LSB)

The DEFAULT RBC EXTENT specifies the extent ID to which RBC READ (10), READ CAPACITY, VERIFY (10) and WRITE (10) commands are mapped. If DEFAULT RBC EXTENT is zero, there is no default extent for RBC commands. Otherwise, the behavior of the data transfer commands is identical to the corresponding extent-relative commands with an EXTENT ID field equal to DEFAULT RBC EXTENT while a READ CAPACITY command returns capacity data for the extent identified by DEFAULT RBC EXTENT.

The DATA LENGTH field specifies the length in bytes of the following data that is available to be transferred. The value of DATA LENGTH does not include the number of bytes in the DATA LENGTH field itself.

NOTE 1 - The total size of the available extent directory may be determined by issuing an EXTENT DIRECTORY command whose EXTENT DIRECTORY OFFSET is zero and whose ALLOCATION LENGTH is greater than or equal to four. The value of DATA LENGTH in the returned information contains the total size of the extent directory.

The EXTENT DIRECTORY is a bit map that specifies whether or not particular extents are active and potentially accessible. The value of EXTENT DIRECTORY OFFSET in the EXTENT DIRECTORY command determines which extents are represented in the bit map. The most significant bit represents extent ID $8 * \text{EXTENT DIRECTORY OFFSET}$. Subsequent, less significant bits represent sequentially increasing extent IDs. Byte n within the extent directory data, where n is greater than or equal to four, represents a range of eight extents from extent ID $(8 * \text{EXTENT DIRECTORY OFFSET}) + ((n - 4) \text{ modulus } 8)$ to extent ID $(8 * \text{EXTENT DIRECTORY OFFSET}) + (((n - 4) \text{ modulus } 8) + 7)$, in order from the byte's most to least significant bit. If a bit is zero, its corresponding extent is not allocated. Otherwise, the corresponding extent is allocated and may be accessible by extent-relative data transfer commands. The QUERY EXTENT command may be used to obtain more information about the extent's size and other characteristics.

5.0b EXTENT MANAGEMENT command

The EXTENT MANAGEMENT command (see Table 2c) provides a means for the initiator to create, delete or in other ways manage extents on the device media. An extent is a logically contiguous range of blocks identified by an EXTENT ID.

Table 2c - EXTENT MANAGEMENT Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0	
0	OPERATION CODE (XXh)								
1	Reserved				ACTION				
2	(MSB)								
3									
4		EXTENT ID							
5									(LSB)
6	(MSB)								
7		DATA FORMAT							(LSB)
8	(MSB)								
9									
10									
11		EXTENT SIZE							
12									
13									(LSB)
14	Reserved								
15	CONTROL = 00h								

The ACTION field specifies the EXTENT MANAGEMENT action requested, as encoded by Table 2d.

Table 2d - EXTENT MANAGEMENT action codes

Code	Name	Description
0	CREATE EXTENT	Create an extent with the data format and size specified.
1	DELETE EXTENT	Delete the extent identified by EXTENT ID and release its medium allocation.
2	LOCK EXTENT	Lock the extent identified by EXTENT ID for exclusive access by RBC commands.
3	UNLOCK EXTENT	Unlock the extent identified by EXTENT ID and release exclusive access.
4	DEFAULT RBC EXTENT	Set the default RBC extent to that specified by EXTENT ID.
5 - 15		Reserved

When the ACTION field specifies CREATE EXTENT, the contents of the EXTENT ID field are unspecified. Otherwise it shall contain an extent ID previously obtained by means of an EXTENT MANAGEMENT command with an ACTION of CREATE EXTENT.

The usage of the DATA FORMAT field is to be determined by the working group.

NOTE A value of zero should probably indicate an unspecified format, that is, there is no foreign protocol (such as AV/C) that can independently access and make sense of the data in this format. How many recording formats are there for AV/C? If vendor-dependent formats are permitted, how are they to be distinguished from each other?

The contents of the EXTENT SIZE field are specified only when the ACTION field equals CREATE EXTENT, in which case the device is requested to create an extent with at least as many blocks as specified by EXTENT SIZE.

When the EXTENT MANAGEMENT command action is CREATE EXTENT, the ASSIGNED EXTENT ID (see Table 2e) shall be returned to the initiator prior to the return of GOOD status for the command.

Table 2e - ASSIGNED EXTENT ID

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1								
2	ASSIGNED EXTENT ID							
3								(LSB)

The ASSIGNED EXTENT ID field specifies the extent ID assigned by the device when an extent is successfully created. The initial state of the extent shall be locked for exclusive RBC access.

5.1a QUERY EXTENT command

The QUERY EXTENT command (see Table 3a) provides a means for the initiator to request information about a particular extent managed by the RBC device.

Table 3a - QUERY EXTENT Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0								
0	OPERATION CODE (XXh)															
1	(MSB)															
2	EXTENT ID															
3																
4									(LSB)							
5																
6	Reserved															
7	(MSB)															
8	ALLOCATION LENGTH															
9	(LSB)															
9	CONTROL = 00h															

The EXTENT ID field identifies the extent for which extent descriptor data is requested.

EXTENT DESCRIPTOR data (see Table 3b) shall be returned to the initiator prior to the return of GOOD status for the command.

Table 3b - EXTENT DESCRIPTOR data

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1								
2	EXTENT ID							
3								(LSB)
4	(MSB)							
5	DATA FORMAT							
6								(LSB)
6	(MSB)							
7								
8								
9	EXTENT SIZE							
10								
11								(LSB)
12	(MSB)							
13	DATA LENGTH							
14								(LSB)
14								
n	EXTENT NAME							

The EXTENT ID field identifies the extent described by the extent descriptor data.

The usage of the DATA FORMAT field is to be determined by the working group.

The contents of the EXTENT SIZE field are specified only when the ACTION field equals CREATE EXTENT, in which case the device is requested to create an extent with at least as many blocks as specified by EXTENT SIZE.

The DATA LENGTH field specifies the length in bytes of the following data that is available to be transferred. The value of DATA LENGTH does not include the number of bytes in the DATA LENGTH field itself.

The EXTENT NAME field is a variable-length field whose usage and meaning is determined by data format. It permits the name of an extent visible *via* another protocol to be correlated with the extent ID used by RBC-2 commands.

5.3 READ CAPACITY command

The READ CAPACITY command (see table 5) provides a means for the initiator to request [either](#) the current capacity of the [entire](#) RBC device [or, if the device supports extent management, the capacity of the current default RBC extent.](#)

Table 5 - READ CAPACITY Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (25h)							
1	Reserved				FREE	TOTAL	LONGLBA	Reserved
2	Reserved							
3								
4								
5								
6								
7								
8								
9	CONTROL = 00h							

[The FREE bit, when one, requests the device to report the capacity of its unallocated addressable media. When the FREE bit is one the TOTAL bit shall be ignored.](#)

[The TOTAL bit specifies the nature of the READ CAPACITY data requested by the initiator. If the device does not support extent management the TOTAL bit shall be ignored and the returned capacity data shall describe all of the device's addressable media. Otherwise, for devices that support extent management, a zero TOTAL bit specifies that the device shall return capacity data for the default RBC extent while a TOTAL bit equal to one specifies that the device shall return the total capacity of the device's addressable media. In the latter case, this shall include addressable media not allocated to any extent.](#)

[When the LONGLBA bit is zero, the device shall return READ CAPACITY data in the format specified by Table 6, otherwise it shall return the data in the format specified by Table 6a.](#)

READ CAPACITY data (see Table 6 [or Table 6a](#)) shall be returned to the initiator prior to sending GOOD status for the command.

Table 6 - READ CAPACITY data (short format)

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1	RETURNED BLOCK ADDRESS							
2								
3								
4	(MSB)							
5	BLOCK LENGTH IN BYTES							
6								
7								

If the LONGLBA bit is zero and the number of blocks exceeds the maximum value that may be specified in the RETURNED BLOCK ADDRESS field, the device server shall set the RETURNED BLOCK ADDRESS field to FFFFFFFFh. The initiator may issue a subsequent READ CAPACITY command whose LONGLBA bit is one.

Table 6a - READ CAPACITY data (long format)

<u>Bit</u> <u>Byte</u>	<u>7</u>	<u>6</u>	<u>5</u>	<u>4</u>	<u>3</u>	<u>2</u>	<u>1</u>	<u>0</u>
<u>0</u>	<u>(MSB)</u>							
<u>1</u>	<u>RETURNED BLOCK ADDRESS</u>							
<u>2</u>								
<u>3</u>								
<u>4</u>								
<u>5</u>								
<u>6</u>								
<u>7</u>	<u>(LSB)</u>							
<u>8</u>	<u>(MSB)</u>							
<u>9</u>	<u>BLOCK LENGTH IN BYTES</u>							
<u>10</u>								
<u>11</u>								

For devices that support extent management, if the FREE bit in the READ CAPACITY command is one, the RETURNED BLOCK ADDRESS represents the device's unallocated medium capacity (*i.e.*, the total amount of addressable medium that may be allocated to one or more extents by the CREATE

EXTENT function of the EXTENT MANAGEMENT command) and the size of the unallocated blocs is specified by BLOCK LENGTH IN BYTES.

For devices that do not support extent management or if the TOTAL bit in the READ CAPACITY command is one, the BLOCK LENGTH IN BYTES and the RETURNED BLOCK ADDRESS are those of the last logical block of the media present in the device. Otherwise, the BLOCK LENGTH IN BYTES and the RETURNED BLOCK ADDRESS are those of the last logical block of the default RBC extent for the media present in the device.

If the device does not contain media, then it shall return status of CHECK CONDITION (02h), sense key of NOT READY (02h), and an ASC of LOGICAL UNIT NOT READY (04h). The ASCQ shall reflect the current state of the device or media.

5.3a READ EXTENT-RELATIVE Command

The READ EXTENT-RELATIVE command (see Table 6b) requests that the device transfer data to the initiator. The most recent data value written in the addressed block within the extent shall be returned.

Table 6b – READ EXTENT-RELATIVE Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0	
0	OPERATION CODE (XXh)								
1	Reserved								
2	(MSB)								
3									
4		RELATIVE BLOCK ADDRESS							
5									
6									
7									(LSB)
8	Reserved								
9	(MSB)	TRANSFER LENGTH							
10									(LSB)
11	(MSB)	EXTENT ID							
12									
13									
14									(LSB)
15	CONTROL = 00h								

The RELATIVE BLOCK ADDRESS field specifies the first block, within the extent, of the range of logical blocks that shall be read.

The TRANSFER LENGTH field specifies the number of contiguous blocks of data that shall be transferred. A TRANSFER LENGTH of zero indicates that no blocks shall be transferred. This condition shall not be considered an error. Any other value indicates the number of blocks that shall be transferred.

The EXTENT ID field specifies the extent upon which the RELATIVE BLOCK ADDRESS is based.

5.6a WRITE EXTENT-RELATIVE Command

The WRITE EXTENT-RELATIVE command (see Table 11a) requests that the device transfer data from the initiator and write it on the medium.

Table 11a – WRITE EXTENT-RELATIVE Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (XXh)							
1	Reserved			FUA		Reserved		
2	(MSB)							
3								
4	RELATIVE BLOCK ADDRESS							
5								
6								
7	(LSB)							
8	Reserved							
9	(MSB)							
10	TRANSFER LENGTH							
11	(MSB)							
12								
13	EXTENT ID							
14	(LSB)							
15	CONTROL = 00h							

A FORCE UNIT ACCESS (FUA) bit of zero indicates that the device may satisfy the command by accessing the cache memory if the WCD bit in RBC mode page 06h (see 5.8.3) is set to zero. For write operations, logical blocks may be transferred directly to the cache memory. GOOD status may be returned to the initiator prior to writing the blocks to the medium. Any error that occurs after GOOD status is returned is a deferred error.

A FUA bit of one indicates that the device shall access the media in performing the command prior to returning GOOD status. A WRITE command shall not return GOOD status until the blocks have actually been written on the media (i.e. the data is not write cached).

If the device supports write caching, FUA support shall be implemented. If write caching is NOT supported then the FUA bit may be ignored.

The RELATIVE BLOCK ADDRESS field specifies the first block, within the extent, of the contiguous range of blocks that shall be written.

The TRANSFER LENGTH field specifies the number of contiguous blocks of data that shall be transferred. A TRANSFER LENGTH of zero indicates that no blocks shall be transferred. This condition shall not be considered an error. Any other value indicates the number of blocks that shall be transferred.

The EXTENT ID field specifies the extent upon which the RELATIVE BLOCK ADDRESS is based.

5.7a VERIFY EXTENT-RELATIVE Command

The VERIFY EXTENT-RELATIVE command (see Table 12a) requests that the device verify the data written on the medium.

Table 12a – VERIFY EXTENT-RELATIVE Command Descriptor Block

Bit Byte	7	6	5	4	3	2	1	0	
0	OPERATION CODE (XXh)								
1	Reserved								
2	(MSB)								
3									
4									
5		RELATIVE BLOCK ADDRESS							
6									
7									(LSB)
8	Reserved								
9	(MSB)								
10		VERIFICATION LENGTH							(LSB)
11	(MSB)								
12									
13		EXTENT ID							
14									(LSB)
15	CONTROL = 00h								

The RELATIVE BLOCK ADDRESS field specifies the first block, within the extent, of the contiguous range of blocks that shall be verified.

The VERIFICATION LENGTH field specifies the number of contiguous blocks of data that shall be verified. A VERIFICATION LENGTH of zero indicates that no blocks shall be verified. This condition shall not be considered an error. Any other value indicates the number of blocks that shall be verified.

The EXTENT ID field specifies the extent upon which the RELATIVE BLOCK ADDRESS is based.

The VERIFY command verifies that the data previously written on the medium is readable without any uncorrectable errors at the time of execution of the command. It does not guarantee the information is complete or valid.