Mode/ Log/ VPD Pages For Describing Solid State Storage (Revision 3.0 Draft 4)

Martin R. Furuhjelm
Pat LaVarre

Contents

A. Scope

Hosts have a need to tune their behavior to fit the significant differences exhibited by the wide variety of flash memory technologies now sold to the mass market. These needs arise across a broad spectrum of “peripheral device type”s, including h 00 04 05 07 0E (i.e., DASD SBC/ RBC, C/DVD MMC, etc.).

Before this proposal, hosts had to run benchmarks to discover such characteristics, specifically benchmarks new and slow and complex enough to defeat benchmark-detection firmware in the device.

This proposal gives the device a standard way to describe its own solid storage media and also gives the host a standard place to store its own estimate of the storage capacity and health of that media.

The host may discover that a device supports this proposed feature by testing for the existence of the new mode, log, or VPD pages.

B. Conventions

In this proposal, we provide written Scope, Conventions, and the *italicized comments* as informative rationale to help accelerate discussion of this proposal.

We expect the committee will choose to delete this non-normative rationale before agreeing to add a version of this proposal into SCSI as an enhancement. Also we have not yet chosen the TBD page codes, section numbers, and table numbers.

As further rationale, we hope to add cross-references to where SCSI already defines the rules that govern when the content of these pages may change.

C. Solid State Storage mode page

We designed this mode page to replace the October 04-362r0.pdf Data Out to "BYTE 18" of the mode page there, described then as: "Low Nibble; Media Health Status; High Nibble; Storage Capacity Status".

7 Parameters for all device types

7.4 Mode parameters

7.4.n Solid State Storage mode page

The Solid State Storage mode page (see table NNN) gives the application client a standard place to store its own estimate of the storage capacity and health of the media. The logical unit may display these estimates on the device server itself.

A device that makes the solid state storage mode page available shall also make the solid state storage log page and the solid state storage VPD page available.

Table NNN – Solid State Storage mode page

Bit Byte	7	6	5	4	3	2	1	0
0	PS(1b)	SPF(0b)	PAGE CODE (TBD)					
1	*ADDITIONAL* PAGE LENGTH (0Fh)							
2-9	Reserved							
10	HOST MEDIA STORAGE CAPACITY							
11	HOST MEDIA HEALTH							
12-15	HOST SIGNATURE							

Each field of this page appears at a byte offset aligned to its byte length.

The PS bit, SPF bit, PAGE CODE field, and PAGE LENGTH field are described in 7.4.5 *Mode page and subpage formats and page codes*.

The logical unit shall set PS to 1b, the SPF to 0b, and the PAGE LENGTH to 06h, as the table shows. The logical unit may, but should not, implement a distinction between current and saved mode pages.

("Implements no distinction" is a quote from SPC "6.7" "MODE SELECT(6) command" for how we ask for a Mode Select of this page with SP clear to die via INVALID FIELD IN CDB - ILLEGAL REQUEST - CHECK CONDITION.)

The HOST MEDIA STORAGE CAPACITY field specifies the percent of free space remaining.

Table NNN – HOST MEDIA STORAGE CAPACITY field

Code		Definition
00h	(0d)	No Information Provided
01h – 64h	(1d – 100d)	Percent Free, Initially 100%
65h – FEh	(2d – 254d)	Reserved
FFh	(255d)	Full = 0% Calculated

The HOST MEDIA HEALTH field specifies the percent of media life remaining.

Table NNN – HOST MEDIA HEALTH field

Code		Definition
00h	(0d)	No Information Provided
01h – 64h	(1d – 100d)	Percent Good, Initially 100%
65h – FEh	(101d – 254d)	Reserved
FFh	(255d)	End of Life = 0% Calculated

The HOST SIGNATURE field specifies another four bytes for the device server to save with the page. The application client may use these bits to more confidently decide if a compatible client calculated the other fields of this page.

D. Solid State Storage log page

We designed this log page to replace the October 04-362r0.pdf Data In from "BYTE 18" of the mode page there, described then as: "Low Nibble; Media Health Status; High Nibble; Storage Capacity Status".

7. Parameters for all device types

7.2 Log parameters

7.2.n Solid State Storage log page

The Solid State Storage log page (see table NNN) provides the logical unit a standard place to report its own estimate of the storage capacity and health of the media. The application client may display these estimates.

A device that makes the solid state storage log page available shall also make the solid state storage mode page and the solid state storage VPD page available.

Table NNN – Solid State Storage log page

Bit Byte	7	6	5	4	3	2	1	0
0	PAGE CODE (TBD)							
1	Reserved							
2-3	(MSB) … *ADDITIONAL* PAGE LENGTH (0Ch) … (LSB)							
4	(MSB) … PARAMETER CODE (TBD) Solid State Storage … (LSB)							
5	DU	DS	TSD	ETC	TMC		LBIN	LP
6	*ADDITIONAL* PARAMETER LENGTH (0Ah)							
7-8	Reserved							
10	DEVICE MEDIA STORAGE CAPACITY							
11	DEVICE MEDIA HEALTH							
12-15	DEVICE SIGNATURE							

Each field of this page appears at a byte offset aligned to its byte length.

The DU DS TSD ETC TMC LBIN LP bits and fields, collectively referred to as the PARAMETER CONTROL byte, are described by spc3r21b.pdf Table 193 *Log parameter*.

The DEVICE MEDIA STORAGE CAPACITY field contains the percent of free space remaining.

Table NNN – DEVICE MEDIA STORAGE CAPACITY field

Code		Definition
00h	(0d)	No Information Provided
01h – 64h	(1d – 100d)	Percent Free, Initially 100%
65h – FEh	(2d – 254d)	Reserved
FFh	(255d)	Full = 0% Calculated

The DEVICE MEDIA HEALTH field contains the percent of media life remaining.

Table NNN – DEVICE MEDIA HEALTH field

Code		Definition
00h	(0d)	No Information Provided
01h – 64h	(1d – 100d)	Percent Good, Initially 100%
65h – FEh	(101d – 254d)	Reserved
FFh	(255d)	End of Life = 0% Calculated

The DEVICE SIGNATURE field contains another four bytes to distinguish the algorithms the logical unit applied to calculate the other fields of this page. The application client may use these bits to more confidently decide if a compatible logical unit calculated the other fields of this page.

E. Solid State Storage VPD page

We designed this VPD page to replace the rest of the October 04-362r0 Data In from the mode page there.

7 Parameters for all device types

7.6 Vital product data parameters

7.6.n Solid State Storage VPD page

The Solid State Storage VPD page (see table NNN) describes the structure of the media in enough detail to let the application client adapt to the significant differences exhibited by a wide variety of solid state storage technologies.

A device that makes the solid state storage VPD page available shall also make the solid state storage mode page and the solid state storage log page available.

Table NNN – Solid State Storage VPD page

Bit Byte	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (TBD)							
2	Reserved							
3	*ADDITIONAL* PAGE LENGTH (NNh)							
TBD	Reserved							
TBD	Reserved						WC	WT
TBD	Reserved						PSI	BB
TBD	MVOL							
TBD	SOLID STATE MEDIA TYPE							
TBD	MIN WRITES PER ERASE *(exponential scale)*							
TBD	MAX PARTIAL WRITES (3:0)							
TBD	ECC DETECTION SPAN							
TBD	ECC CORRECTION SPAN							
TBD	Reserved							
TBD	MIN SUSTAINED SEQUENTIAL READ SPEED (7:0)							
TBD	MIN SUSTAINED SEQUENTIAL WRITE SPEED (7:0)							
TBD	MAX RANDOM READ ACCESS *(exponential scale)*							

TBD	MAX RANDOM WRITE ACCESS *(exponential scale)*
TBD	Reserved
TBD	BITS PER CELL (1:0)
TBD	BYTES PER SECTOR (1:0)
TBD	SECTORS PER PAGE (1:0)
TBD	PAGES PER ERASE (3:0)
TBD	ERASE BLOCKS PER DIE (7:0)
TBD	DIE WIDTH *(exponential scale)*
TBD	DIE COUNT (1:0)
TBD	Reserved
TBD	JEDEC MANUFACTURER ID (7:0)
TBD	JEDEC PRODUCT ID (7:0)

Each field of this page appears at a byte offset aligned to its byte length.

The PERIPHERAL QUALIFIER field and the PERIPHERAL DEVICE TYPE field are defined in SPC-3.

The PAGE CODE field shall be set to (TBD).

The *ADDITIONAL* PAGE LENGTH field is defined in SPC-3.

7.6.n.1 Medium Volatility

A WC bit set to one indicates the logical unit supports FUA.

A WT bit set to one indicates the logical unit includes a write cache and supports controlling that cache by mode pages 6 and/or 8.

A PSI bit set to one indicates the logical unit conforms to the 03-388r3 specification for Power Supply information.

A BB bit set to one indicates the logical unit supplies a battery backup in conformance with 03-388r3 SPC-3, SBC-2 Nonvolatile Caches.

The MVOL field is defined in table NNN.

0h	-	No Information Provided
1h	-	Not Changeable (ROM)
2h	-	Non-Volatile (e.g., flash memory)

3h	-	Volatile – (RAM)
4h-FFh	-	Reserved

The MVOL field exists to divide the present & future possibilities for SOLID STATE MEDIA TYPE into less specific classes.

The SOLID STATE MEDIA TYPE field is defined in table NNN.

0h	-	No Information Provided
1h	-	ROM (Read Only)
2h	-	OTP (One Time Programmable)
3h	-	NOR Flash
4h	-	NAND Flash
5h	-	AND Flash
6h	-	AG-AND Flash
7h	-	Floating Gate Flash
8h-FFh	-	Reserved

The MIN WRITES PER ERASE field contains zero if no information provided, else the minimum number of program cycles that the medium supports before requiring an erase. The MIN WRITES PER ERASE field is defined on an exponential scale in table NNN.

0h	-	No Information Provided
1h	-	No erase (e.g., OTP, WORM)
2h	-	10 or less
3h	-	100
4h	-	1,000
5h	-	10,000
6h	-	100,000 *(ordinary flash, 2004)*
7h	-	1,000,000
8h	-	10,000,000
9h	-	100,000,000
Ah	-	1,000,000,000
Bh-FEh	-	Reserved
FFh	-	Unlimited

A partial write is a write of less than all the sectors of a page. Some Solid State Media (e.g., flash) allows for a limited number of partial writes in a page before the page has to be erased in order to continue writing to it.

The MAX PARTIAL WRITES (3:0) field contains zero if no information provided, else the number of partial writes that the solid state medium allows before requiring an erase, encoded as defined in Table NNN.

0h	-	No Information Provided
1h	-	Write Once Only
2h – FEh		1d - 254d Partial writes allowed before erase
FFh – FFFF:FFFEh		Reserved
FFFF:FFFFh	-	Unlimited partial writes allowed (erase still needed to rewrite)

The ECC DETECTION SPAN field contains zero if no information provided, else the number of bits in error that the ECC code can detect. *Detecting two wrong bits was normal for flash in 2004.*

The ECC CORRECTION SPAN field contains zero if no information provided, else the number of bits in error that the ECC code can correct. *Correcting one wrong bit was normal for flash in 2004.*

7.6.n.2 Medium Speed

The MIN SUSTAINED SEQUENTIAL READ SPEED (7:0) field contains zero if no information provided, else the min sustained sequential read speed, in units of bytes per second.

The MIN SUSTAINED SEQUENTIAL WRITE SPEED (7:0) field contains zero if no information provided, else the min sustained sequential read speed, in units of bytes per second.

The MAX RANDOM READ ACCESS field is defined on an exponential scale in table NNN.

.

0h	-	No Information Provided
1h	-	10s or more
2h	-	1000 ms
3h	-	100 ms
4h	-	10 ms
5h	-	1 ms
6h	-	100 us
7h	-	10 us
8h	-	1 us *(ordinary flash, 2004)*
9h	-	100 ns
Ah	-	10 ns

Bh	-	1 ns
Ch	-	100 ps
Dh	-	10 ps
Eh	-	1ps
Fh-FFh	-	Reserved

The MAX RANDOM WRITE ACCESS field is defined by the same exponential scale as for reading, in table NNN above.

7.6.n.3 Medium Structure

The following fields describe a storage medium structure which divides the medium into chip dies, which divide into erase blocks, which divide into pages, which divide into sectors of 512 bytes each. The medium may group bits together into cells.

The BITS PER CELL (1:0) field contains zero if no information provided, else the bits per cell. *Less than five bits per cell was normal for flash in 2004.*

The BYTES PER SECTOR (1:0) field contains zero if no information provided, else the number of bytes per sector, i.e., 512.

The SECTORS PER PAGE (1:0) field contains zero if no information provided, else the number of sectors per page. *"Page Size" is the name that many solid state storage data sheets give to this value. 4 sectors per page was normal for "large block" flash, 1 sector per page was normal for "small block" flash in 2004.*

The PAGES PER ERASE (3:0) field contains zero if no information provided, else the number of pages per erase block. *"Block Size" is the name that many solid state storage data sheets give to this value. 128 KiB (often 256 sectors divided into 64 pages) per erase block was normal for flash in 2004.*

The ERASE BLOCKS PER DIE (7:0) field contains zero if no information provided, else the number of erase blocks per chip die. *1 or 2 GiB per die was normal for flash in 2004.*

The DIE COUNT (1:0) field contains zero if no information provided, else the number of chip dies. *1 or 2 or 4 was normal for flash in 2004.*

The DIE WIDTH field contains zero if no information provided, else the bits per chip die, defined on an exponential scale in table NNN.

0h	-	No Information Provided
1h	-	1
2h	-	2

3h	-	4
4h	-	8
5h	-	16
6h	-	32
7h	-	64 *(64 or less normal for flash in 2004)*
8h	-	128
9h	-	256
Ah	-	512
Bh	-	1024
Ch-FFh	-	Reserved

7.6.n.4 Medium Identity

The JEDEC MANUFACTURER ID (7:0) field contains all nulls (00h) to indicate no information provided, else N bytes of code from the list that JEDEC maintains as Publication JEP106 "Standard Manufacturer's Identification Code", followed by 8 – N bytes of null padding. *For example, the bytes h 7F:7F:7F:7F 83:00:00:00 00:00:00:00 00:00:00:00 represent "Teknovus", listed as JEDEC manufacturer 03H with odd parity 80H of the bank found past four 7FH continuations.*

The JEDEC PRODUCT ID (7:0) field contains all nulls (00h) to indicate no information provided, else bits defined by the JEDEC MANUFACTURER to distinguish one solid state storage product from another, null-padded on the right (i.e., highest offsets).

Change Log

In reverse chronological order, with reference to http://t10.org/doc04.htm

4. TBD

3. Tuesday, January 18, 2005

04-362r2 Mode/ Log/ VPD Pages For Describing Solid State Storage, Pat LaVarre, PDF (maybe less than 40,000 bytes)

VPD fields merged in from 04-362r0.

2. Monday, January 17, 2005

04-362r1 SBC-2: Solid State Storage Description Mode/ Log/ VPD Description, Martin R. Furuhjelm, PDF (34,597 bytes)

SCSI in general substituted for SCSI thru USB in particular.

Data divided into mode, log, and VPD pages.

Scope & Conventions summarized.

1. October 2004

04-362r0 SBC-2: Proposal for USB Solid State Drive Mode Sense specification, Martin R. Furuhjelm, PDF (30831 bytes).

First proposal published.

Revisions recommended by 04-367r1 Minutes of CAP Working Group - Nov 9-10, 2004 Weber & Lohmeyer HTM (18323)