

T10/08-386 revision 1

Date: 11/3/08

To: T10 Committee (SCSI)

From: George Penokie (LSI)

Subject: SPC-4: Cache hits and power on statistics

1 Overview

There are some performance statistics relating to cache hits that would be useful for applications that are not included in any of the statistics log pages. This proposal adds in the following statistics into a Cache Statistics and Performance log page:

- a) read cache hits;
- b) writes from cache, and
- c) write cache hits.

Also, this proposal adds a power on timer into the new Cache Statistics and Performance log page as there is none that is currently defined that uses the time interval descriptor and only counts time since the last hard reset event.

In addition this proposal fixes the overview that was not modified to reflect the addition of the FUA statistics.

1.0.1 Log page codes for all device types

The page code assignments for the log pages are listed in table 1.

Table 1 — Log page codes

Page Code ^a	Subpage Code ^a	Log Page Name	Reference
0Fh	00h	Application Client	7.2.3
01h	00h	Buffer Over-Run/Under-Run	7.2.4
19h	00h	General Statistics and Performance	7.2.13
19h	01h to 1Fh	Group Statistics and Performance (1 to 31)	7.2.13
19h	20h	Cache Statistics and Performance	1.0.2.3
2Fh	00h	Informational Exceptions	7.2.6
0Bh	00h	Last <i>n</i> Deferred Errors or Asynchronous Events	7.2.7
07h	00h	Last <i>n</i> Error Events	7.2.8
06h	00h	Non-Medium Error	7.2.9
18h	00h to FEh	Protocol Specific Port ^b	7.2.10
03h	00h	Read Error Counter	7.2.5
04h	00h	Read Reverse Error Counter	7.2.5
10h	00h	Self-Test Results	7.2.11
0Eh	00h	Start-Stop Cycle Counter	7.2.12
00h	00h	Supported Log Pages	7.2.14
00h	FFh	Supported Log Pages and Subpages	7.2.15
01h to 3Fh	FFh	Supported Subpages	7.2.16
0Dh	00h	Temperature	7.2.17
05h	00h	Verify Error Counter	7.2.5
02h	00h	Write Error Counter	7.2.5
08h to 0Ah	00h to FEh	Reserved (may be used by specific device types)	
0Ch	00h to FEh	Reserved (may be used by specific device types)	
11h to 17h	00h to FEh	Reserved (may be used by specific device types)	
19h	20h 1h to FEh	Reserved	
1Ah to 2Eh	00h to FEh	Reserved (may be used by specific device types)	
3Fh	00h to FEh	Reserved	
30h to 3Eh	00h to FEh	Vendor specific	
Annex D contains a listing of log pages codes and subpage codes in numeric order.			
^a All page code and subpage code combinations not shown in this table are reserved.			
^b Each SCSI transport protocol standard (see 3.1.140) may define a different name for these log pages.			

....

1.0.2 Statistics and Performance log pages

1.0.2.1 Statistics and Performance log pages overview

The Statistics and Performance log pages consist of a General Statistics and Performance log page ~~and~~, up to 31 Group Statistics and Performance log pages, ~~and a Cache Memory Statistics and Performance log page~~. Each Group Statistics and Performance log page ~~s~~ only collects statistics and performance information for the group number specified in a read CDB or a write CDB.

The General Statistics and Performance log page (see 1.0.2.2) provides the following statistics and performance results associated to the addressed logical unit:

- a) Statistics and Performance log parameters:
 - A) Number of read commands;
 - B) Number of write commands;
 - C) Number of read logical blocks transmitted by a target port;
 - D) Number of write logical blocks received by a target port;
 - E) Read command processing time;
 - F) Write command processing time;
 - G) Sum of the command weights of the read commands plus write commands; and
 - H) Sum of the weighted command time of the read commands plus write commands;
- b) Idle Time log parameter:
 - A) Idle time; ~~and~~
- c) Time Interval log parameter:
 - A) Time interval; ~~and~~
- d) Force Unit Access Statistics and Performance log parameter:
 - A) Number of read FUA commands;
 - B) Number of write FUA commands;
 - C) Number of read FUA_NV commands;
 - D) Number of write FUA_NV commands; and
 - E) Read FUA command processing intervals.

The Group Statistics and Performance log pages (see 7.2.13.3) provide the following statistics and performance results associated to the addressed logical unit and the GROUP NUMBER field:

- a) Statistics and Performance log parameters:
 - A) Number of read commands;
 - B) Number of write commands;
 - C) Number of read logical blocks transmitted by a target port;
 - D) Number of write logical blocks received by a target port;
 - E) Read command processing time; and
 - F) Write command processing time; ~~and~~
- b) Force Unit Access Statistics and Performance log parameter:
 - A) Number of read FUA commands;
 - B) Number of write FUA commands;
 - C) Number of read FUA_NV commands;
 - D) Number of write FUA_NV commands; and
 - E) Read FUA command processing intervals.

The Cache Memory Statistics and Performance log page (see 1.0.2.3) provides the following statistics and performance results associated with the addressed logical unit:

- a) Number of read commands received through a SCSI target port that resulted in user data being read from a cache memory;
- b) Number of write commands received through a SCSI target port that resulted in user data being written to cache memory that did not result in a write of any user data to media before the write command being processed is completed;
- c) Number of write operations initiated to move user data from cache memory to media; and
- d) Time from last hard reset (see SAM-4).

In the Statistics and Performance log pages, read and write commands are those shown in table 2.

Table 2 — Statistics and Performance log pages read and write commands

Read commands ^a	Write commands ^a
READ(6)	WRITE(6)
READ(10)	WRITE(10)
READ(12)	WRITE(12)
READ(16)	WRITE(16)
READ(32)	WRITE(32)
	WRITE AND VERIFY(10)
	WRITE AND VERIFY(12)
	WRITE AND VERIFY(16)
	WRITE AND VERIFY(32)
^a See SBC-3.	

1.0.2.2 General Statistics and Performance log page

....

1.0.2.3 Cache Memory Statistics and Performance log page (this section is all new)

The Cache Memory Statistics and Performance log page (see table 3) provides logging of statistics and performance of read and write operations using cache memory.

Table 3 — Cache Memory Statistics and Performance log page

Bit Byte	7	6	5	4	3	2	1	0
0	DS	SPF (1b)	PAGE CODE (19h)					
1	SUBPAGE CODE (20h)							
2	(MSB)	PAGE LENGTH (3428h)						(LSB)
3								
Cache Memory Statistics and Performance log parameters								
4	Cache Memory log parameter (required) (see table							
31	4)							
32	Time From Last Hard Reset log parameter							
43	(required) (see table 5)							

The DS bit, SPF bit, PAGE CODE field, and PAGE LENGTH field are described in 7.2.1.

Cache Memory Statistics and Performance log parameters not defined by this standard are reserved.

Table 4 shows the Cache Memory log parameter format.

Table 4 — Cache Memory log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0001h) _____							
								(LSB)
2	DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3	PARAMETER LENGTH (108h)							
4	(MSB) _____							
11	NUMBER OF READ CACHE HITS _____							
								(LSB)
12	(MSB) _____							
19	NUMBER OF WRITE CACHE HITS _____							
								(LSB)
20	(MSB) _____							
27	NUMBER OF WRITES FROM CACHE _____							
								(LSB)

The PARAMETER CODE field set to 0001h identifies the log parameter being transferred as the Cache Memory log parameter.

The FORMAT AND LINKING field ~~for the Cache Memory log parameter in the Statistics and Performance log page~~ shall be set to 10b, indicating that the parameter is a data counter parameter. The values of the bits and fields in the parameter control byte for a data counter parameter are defined in 7.2.1.2.2.2.

The PARAMETER LENGTH field specifies the length in bytes ~~of the statistics and performance fields~~ in the log parameter that follow.

The NUMBER OF READ CACHE HITS field contains the number of read commands (see 1.0.2.1) received by the logical unit that resulted in user data being read from cache memory. that does not result in a read of any of the requested user data from media before the read command completes.

The NUMBER OF WRITE CACHE HITS field contains the number of write commands (see 1.0.2.1) received by the logical unit that resulted in user data being written to cache memory that did not result in a write of any user data to media before the write command completed.

The NUMBER OF WRITES FROM CACHE field contains the number of write operations initiated to move user data from cache memory to media.

Table 5 shows the Time From Last Hard Reset log parameter format.

Table 5 — Time From Last Hard Reset log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0002h) _____							
								(LSB)
2	DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3	PARAMETER LENGTH (08h)							
4	(MSB) _____							
11	LAST HARD RESET INTERVALS _____							
								(LSB)

The PARAMETER CODE field set to 0002h identifies the log parameter being transferred as the Time From Last Hard Reset log parameter.

The FORMAT AND LINKING field ~~for the Time From Last Hard Reset log parameter in the Statistics and Performance log page~~ shall be set to 10b, indicating that the parameter is a data counter parameter. The values of the bits and fields in the parameter control byte for a data counter parameter are defined in 7.2.1.2.2.2.

The PARAMETER LENGTH field ~~specifies~~indicates the length in bytes of the LAST HARD RESET INTERVALS field that follows.

The LAST HARD RESET INTERVALS field contains the number of time intervals that have occurred since a hard reset was processed by the logical unit.

This time is calculated as follows:

$$\text{time} = (\text{time intervals since last hard reset} \times \text{time interval})$$

where:

time intervals since last hard reset is the contents of the LAST HARD RESET INTERVALS field; and

time interval is the value represented in the time interval descriptor of the Time Interval log parameter (see table 6).

Table 6 shows the Time Interval log parameter format.

Table 6 — Time Interval log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0003h) _____ (LSB)							
2	DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3	PARAMETER LENGTH (08h)							
4	Time interval descriptor <u>(see table 7)</u> _____							
11								

Editor's Note 1: This log parameter and its description is identical to what is already described in the general statistics and performance log page and could be deleted with just a reference to that if this stays in SPC.

The PARAMETER CODE field set to 0003h identifies the log parameter being transferred as the Time Interval log parameter.

The FORMAT AND LINKING field ~~for the Time Interval log parameter in the Statistics and Performance log page~~ shall be set to 10b, indicating that the parameter is a data counter parameter. The values of the bits and fields in the parameter control byte for a data counter parameter are defined in 7.2.1.2.2.2.

The PARAMETER LENGTH field specifies the length in bytes of the ~~IDLE TIME INTERVALS field~~time interval descriptor that follows.

The time interval descriptor (see table 7) contains the time interval in seconds.

Table 7 — Time interval descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	EXPONENT						(LSB)
3							(LSB)	
4	(MSB)	INTEGER						(LSB)
7							(LSB)	

The EXPONENT field contains the negative power of 10 exponent to multiply with the INTEGER field (e.g., a value of 9 represents 10^{-9})

When multiplied by the exponent, the INTEGER field contains the value that represents one time interval (e.g., a value of 5 in the INTEGER field and a value of 9 in the EXPONENT field represents a time interval of 5×10^{-9} seconds or 5 nanoseconds).