

T10/08-386 revision 3

Date: 1/14/09

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.

Revision 2: Moved the new sublog page into it's own section, made all the counters into separate log parameters, added a read from medium log parameter, and several other editorial changes were made.

Revision 3: Made it clear what commands are associated with the cache operations and made several other editorial changes.

2 Changes to SPC-4

3.1 Definitions

3.1.xx protection information: Fields appended to each logical block that contain a cyclic redundancy check (CRC), an application tag, and a reference tag. See SBC-3.

3.1.x user data: Data contained in logical blocks that is not protection information (see 3.1.xx).

2.0.1 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 Memory Statistics	2.0.3
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.			

....

2.0.2 Statistics and Performance log pages

2.0.2.1 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. Each Group Statistics and Performance log pages 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 2.0.2.2) 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 parameters:
 - A) Number of read commands with the FUA bit set to one;
 - B) Number of write commands with the FUA bit set to one;
 - C) Number of read commands with the FUA_NV bit set to one;
 - D) Number of write commands with the FUA_NV bit set to one; and
 - E) Read command with the FUA bit set to one 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 parameters:
 - A) Number of read commands with the FUA bit set to one;
 - B) Number of write commands with the FUA bit set to one;
 - C) Number of read commands with the FUA_NV bit set to one;
 - D) Number of write commands with the FUA_NV bit set to one; and
 - E) Read command with the FUA bit set to one processing intervals.

In the Statistics and Performance log pages, read [commands](#) 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.	

2.0.2.2 General Statistics and Performance log page

....

2.0.3 [Cache Memory Statistics log page](#)

2.0.3.1 [Cache Memory Statistics log page overview](#)

[The Cache Memory Statistics log page provides the following statistics and performance results associated with the addressed logical unit:](#)

- a) [Cache Memory Statistics log parameters:](#)
 - A) [Number of read cache memory hits:](#)
 - B) [Number of reads to cache memory:](#)
 - C) [Number of write cache memory hits: and](#)
 - D) [Number of writes from cache memory:](#)
- b) [Hard Reset log parameter:](#)
 - A) [Time from last hard reset \(see SAM-4\).](#)
[and](#)
- c) [Time Interval log parameter:](#)
 - A) [Time interval:](#)

In the Cache Memory Statistics pages, read cache operations and write cache operations shall only result from the read commands and write commands shown in table 3.

Table 3 — Cache Memory Statistics log pages cache commands

Read commands ^a	Write commands ^a
READ(6)	SYNCHRONIZE CACHE (10)
READ(10)	SYCNHRONIZE CACHE (16)
READ(12)	WRITE(6)
READ(16)	WRITE(10)
READ(32)	WRITE(12)
PRE-FETCH (10)	WRITE(16)
PRE-FETCH (16)	WRITE(32)
	WRITE AND VERIFY(10)
	WRITE AND VERIFY(12)
	WRITE AND VERIFY(16)
	WRITE AND VERIFY(32)

^a See SBC-3.

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

The Cache Memory Statistics and Performance log page (see table 4) provides logging of statistics and performance of read cache operations and write cache operations using cache memory.

Table 4 — 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 (34h n-3)						(LSB)
3								
Cache Memory Statistics and Performance log parameters								
4	Cache memory statistics log parameter (required)(see table 6)[first]							
34-15	-							
32n-13	-							
43n	Time Interval log parameter (see table 11) Cache memory statistics log parameter [last]							

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 5 defines the parameter codes.

Table 5 — Cache Memory Statistics log page parameter codes

Parameter code	Description
0001h	Read Cache Memory Hits
0002h	Reads To Cache Memory
0003h	Write Cache Memory Hits
0004h	Writes From Cache Memory
0005h	Time From Last Hard Reset
0006h	Time Interval log parameter
all other values	Reserved

Table 6 shows the Read Cache Memory Hits Memory log parameter format.

Table 6 — Read Cache Memory Hits 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 (1008h) _____							
4	(MSB) _____							
11	NUMBER OF READ CACHE MEMORY HITS _____ (LSB)							

The PARAMETER CODE field set to 0001h identifies the log parameter being transferred as the Read Cache Memory Hits 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 MEMORY HITS field contains the number of read commands ~~(see 2.0.3.1)(see table 3)~~ received ~~by the logical unit that resulted in user data being read from cache memory, the medium on an I_T nexus that resulted in read cache operations (see 2.0.3.1) where:~~

- a) user data is read from cache memory; and
- b) no user data is read from the medium before the read command being processed is completed.

The NUMBER OF READ CACHE MEMORY HITS field shall not be modified as a result of any read command that contains an FUA bit set to one or an FUA_NV bit set to one.

The contents of the NUMBER OF READ CACHE MEMORY HITS field shall be set to zero as the result of a hard reset condition (see SAM-4).

Table 7 shows the Reads To Cache Memory log parameter format.

Table 7 — Reads To Cache Memory log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	<u>(MSB)</u>							
1	PARAMETER CODE (0002h)							
	<u>(LSB)</u>							
2	<u>DU</u>	<u>Obsolete</u>	<u>ISD</u>	<u>ETC</u>	<u>TMC</u>	<u>FORMAT AND LINKING</u>		
3	PARAMETER LENGTH (08h)							
4	<u>(MSB)</u>							
	NUMBER OF READS TO CACHE MEMORY							
11	<u>(LSB)</u>							

The PARAMETER CODE field set to 0004h identifies the log parameter being transferred as the Reads To Cache Memory log parameter.

The FORMAT AND LINKING field 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 in the log parameter that follow.

The NUMBER OF READS TO CACHE MEMORY field contains the number of read commands (see table 3) initiated to move user data to cache memory from the medium.

The NUMBER OF READS TO CACHE MEMORY field shall not be modified as a result of any read operation that contains an FUA bit set to one or an FUA_NV bit set to one.

The contents of the NUMBER OF READS TO CACHE MEMORY field shall be set to zero as the result of a hard reset condition (see SAM-4).

Table 8 shows the Write Cache Memory Hits log parameter format.

Table 8 — Write Cache Memory Hits log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	<u>(MSB)</u>							
1	PARAMETER CODE (0003h)							
	<u>(LSB)</u>							
2	<u>DU</u>	<u>Obsolete</u>	<u>ISD</u>	<u>ETC</u>	<u>TMC</u>	<u>FORMAT AND LINKING</u>		
3	PARAMETER LENGTH (08h)							
4	<u>(MSB)</u>							
	NUMBER OF WRITE CACHE MEMORY HITS							
11	<u>(LSB)</u>							

The PARAMETER CODE field set to 0003h identifies the log parameter being transferred as the Write Cache Memory Hits log parameter.

The FORMAT AND LINKING field 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 in the log parameter that follow.

The NUMBER OF WRITE CACHE MEMORY HITS field contains the number of write commands (see table 3) received on an I_T nexus that resulted in write cache operations (see 2.0.3.1) where:

- a) user data is written to cache memory; and
- b) no user data is written to the medium before the write command being processed is completed.

The NUMBER OF WRITE CACHE MEMORY HITS field shall not be modified as a result of any write command that contains an FUA bit set to one or an FUA_NV bit set to one.

The contents of the NUMBER OF WRITE CACHE MEMORY HITS field shall be set to zero as the result of a hard reset condition (see SAM-4).

Table 9 shows the Writes From Cache Memory log parameter format.

Table 9 — Writes From Cache Memory log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0004h)						_____ (LSB)	
2	DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3	PARAMETER LENGTH (08h)							
4	(MSB) _____							
11	NUMBER OF WRITES FROM CACHE MEMORY						_____ (LSB)	

The PARAMETER CODE field set to 0004h identifies the log parameter being transferred as the Writes From Cache Memory log parameter.

The FORMAT AND LINKING field 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 in the log parameter that follow.

The NUMBER OF WRITES FROM CACHE MEMORY field contains the number of write commands (see table 3) initiated to move user data from cache memory to the medium.

The NUMBER OF WRITES FROM CACHE MEMORY field shall not be modified as a result of any write operation that contains an FUA bit set to one or an FUA_NV bit set to one.

The contents of the NUMBER OF WRITES FROM CACHE MEMORY field shall be set to zero as the result of a hard reset condition (see SAM-4).

Table 10 shows the Time From Last Hard Reset log parameter format.

Table 10 — Time From Last Hard Reset log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (00025h) _____ (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 00025h 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 11).

Table 11 shows the Time Interval log parameter format.

Table 11 — Time Interval log parameter

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (00036h) _____ (LSB)							
2	DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3	PARAMETER LENGTH (08h)							
4	_____							
11	Time interval descriptor <u>(see table 12)</u> _____							

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. However, the changes shown here should be incorporated into the current wording.

The PARAMETER CODE field set to 00036h 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 401b, indicating that the parameter is a ~~data-counter~~ list binary 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 12) contains the time interval in seconds.

Table 12 — Time interval descriptor

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