

Date: June 27, 2005

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SPC-4: Statistics and Performance Log Pages

1 Overview

In large networks it is becoming important, and in some cases necessary, to monitor the performance of logical units to determine that the customer is receiving the level of throughput they have contracted. This is best handled by setting up log pages that would return a set of specified performance parameters.

This proposal fills an important gap in the ability to monitor block storage resources. It provides a mechanism, not previously available, by which block devices are able to measure and report their utilization. In addition, it is, currently, the only way that a block storage device has to report the activity of an identified I/O group. (i.e., by group number).

This proposal defines a log page that will return a set of overall performance parameters and a set of subpage log page, for that log page, that would return performance parameters based on group number.

1.0.1 Statistics and Performance subpage logs

1.0.1.1 Statistics and Performance subpage overview

The Statistics and Performance subpage logs consist of a General Statistics and Performance subpage log and up to 31 Group Statistics and Performance subpage logs. Each Group Statistics and Performance subpage log only collects statistics and performance information for the group number specified in a read CDB or a write CDB (see table 2).

The General Statistics and Performance subpage log (see 1.0.1.2) provides the following statistics and performance results associated to the addressed logical unit:

- a) number of read commands;
- b) number of write commands;
- c) number of read blocks sent;
- d) number of write blocks received;
- e) read command processing time;
- f) write command processing time;
- g) number read commands and write commands weighted by priority;
- h) time processing read commands and write commands weighted by priority; and
- i) clock increment in nanoseconds.

The Group Statistics and Performance subpage logs (see 1.0.1.3) provide the following statistics and performance results associated to the addressed logical unit and the GROUP NUMBER field:

- a) number of read commands;
- b) number of write commands;
- c) number of read blocks sent;
- d) number of write blocks received;
- e) read command processing time; and
- f) write command processing time.

In the General Statistics and Performance subpage log and the Group Statistics and Performance subpage logs a read command is one of the following commands:

- a) READ(6) command;
- b) READ(10) command;
- c) READ(12) command;
- d) READ(16) command;
- e) READ(32) command;

- f) XDREAD(10) command; or
- g) XDREAD(32) command.

In the General Statistics and Performance subpage log and the Group Statistics and Performance subpage logs a write command is one of the following commands:

- a) WRITE(6) command;
- b) WRITE(10) command;
- c) WRITE(12) command;
- d) WRITE(16) command;
- e) WRITE(32) command;
- f) WRITE AND VERIFY(10) command;
- g) WRITE AND VERIFY(12) command;
- h) WRITE AND VERIFY(16) command;
- i) WRITE AND VERIFY(32) command;
- j) XDWRITE(10) command; or
- k) XDWRITE(32) command.

In the General Statistics and Performance subpage log the weighted command priority of a command is calculated as follows:

$$\text{command weight} = (360\ 360 / \text{priority of the command})$$

where:

priority of the command is the value of the PRIORITY field in the CDB or if the PRIORITY field is set to zero, then the INITIAL PRIORITY field in the Control Extension mode page (see SPC-3) or the last priority assigned to the I_T_L nexus using the SET PRIORITY command.

In the General Statistics and Performance subpage log the weighted time of a command is calculated as follows:

$$\text{weighted command time} = (\text{time processing the command} \times (360\ 360 / \text{priority of the command})).$$

where:

priority of the command is the value of the PRIORITY field in the CDB or if the PRIORITY field is set to zero, then the INITIAL PRIORITY field in the Control Extension mode page (see SPC-3) or the last priority assigned to the I_T_L nexus using the SET PRIORITY command; and

time processing a command shall be measured from the time the task manager places the command into a task until the device server receives SCSI Transport Protocol Service Confirmation that the status from the command has been sent.

1.0.1.2 General Statistics and Performance subpage log

Table 1 specifies the General Statistics and Performance subpage log parameters.

Table 1 — General Statistics and Performance log subpage

Bit Byte	7	6	5	4	3	2	1	0	
0	PAGE CODE (xxh)								
1	SUBPAGE CODE								
2	(MSB)	PAGE LENGTH (62h)							
3								(LSB)	
General Statistics and Performance sublog parameters									
4	Number of Read Commands sublog parameter								
11									
12	Number of Write Commands sublog parameter								
19									
20	Number of Blocks Received sublog parameter								
27									
28	Number of Blocks Sent sublog parameter								
35									
36	Read Command Processing Time sublog parameter								
43									
44	Write Command Processing Time sublog parameter								
51									
62	Weighed Number of Read Commands plus Write Commands sublog								
59	parameter								
70	Weighed Read Command Processing plus Write Command Processing								
77	sublog parameter								
78	Idle Time sublog parameter								
85									
86	Time Interval sublog parameter								
101									

The PAGE CODE and PAGE LENGTH fields are described in 7.2.1.

The SUBPAGE CODE field is as specified in table 2.

Table 2 — Group Statistics and Performance log subpage codes

Subpage Code	Description	Group number ^{a b}
00h	General Statistics and Performance log subpage	00000b
01h	Group Statistics and Performance log subpage (1)	00001b
02h	Group Statistics and Performance log subpage (2)	00010b
03h	Group Statistics and Performance log subpage (3)	00011b
04h	Group Statistics and Performance log subpage (4)	00100b
05h	Group Statistics and Performance log subpage (5)	00101b
06h	Group Statistics and Performance log subpage (6)	00110b
07h	Group Statistics and Performance log subpage (7)	00111b
08h	Group Statistics and Performance log subpage (8)	01000b
09h	Group Statistics and Performance log subpage (9)	01001b
0Ah	Group Statistics and Performance log subpage (10)	01010b
0Bh	Group Statistics and Performance log subpage (11)	01011b
0Ch	Group Statistics and Performance log subpage (12)	00100b
0Dh	Group Statistics and Performance log subpage (13)	01101b
0Eh	Group Statistics and Performance log subpage (14)	01110b
0Fh	Group Statistics and Performance log subpage (15)	01111b
10h	Group Statistics and Performance log subpage (16)	10000b
11h	Group Statistics and Performance log subpage (17)	10001b
12h	Group Statistics and Performance log subpage (18)	10010b
13h	Group Statistics and Performance log subpage (19)	10011b
14h	Group Statistics and Performance log subpage (20)	10100b
14h	Group Statistics and Performance log subpage (21)	10101b
15h	Group Statistics and Performance log subpage (22)	10110b
17h	Group Statistics and Performance log subpage (23)	10111b
18h	Group Statistics and Performance log subpage (24)	11000b
19h	Group Statistics and Performance log subpage (25)	11001b
1Ah	Group Statistics and Performance log subpage (26)	11010b
1Bh	Group Statistics and Performance log subpage (27)	11011b
1Ch	Group Statistics and Performance log subpage (28)	10100b
1Dh	Group Statistics and Performance log subpage (29)	11101b
1Eh	Group Statistics and Performance log subpage (30)	11110b
1Fh	Group Statistics and Performance log subpage (31)	11111b

^a The GROUP NUMBER field is from the read command CDB or the write command CDB (see SBC-3).

^b The statistics and performance information associated with a group number is collected in the corresponding Group Statistics and Performance log subpage.

Table 3 shows the format of Number of Read Commands sublog parameter.

Table 3 — Number of Read Commands sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0001h) _____ (LSB)							
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	NUMBER OF READ COMMANDS _____ (LSB)							

The PARAMETER CODE field set to 0001h identifies the sublog parameter being transferred as the Number of Read Commands sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

Table 4 — Parameter control bits for self-test results log parameters

Bit	Value	Description
DU	0	Value provided by device server
DS	0	Device server supports saving of parameter
TSD	0	Device server manages saving of parameter
ETC	0	No threshold comparison is made on this value
TMC	xx	Ignored when the ETC bit is set to zero
LBIN	x	Ignored when the LP bit is set to zero
LP	0	The parameter is a list parameter

The PARAMETER LENGTH field shall contain 04h.

The NUMBER OF READ COMMANDS field contains the number of read commands (see 1.0.1.1) received by the addressed logical unit.

Table 5 shows the format of Number of Write Commands sublog parameter.

Table 5 — Number of Write Commands sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0002h) _____ (LSB)							
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	NUMBER OF WRITE COMMANDS _____ (LSB)							

The PARAMETER CODE field set to 0002h identifies the sublog parameter being transferred as the Number of Write Commands sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The NUMBER OF WRITE COMMANDS field contains the number of write commands (see 1.0.1.1) received by the addressed logical unit.

Table 6 shows the format of Number of Blocks Received sublog parameter.

Table 6 — Number of Blocks Received sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0003h)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	NUMBER OF BLOCKS RECEIVED							(LSB)

The PARAMETER CODE field set to 0003h identifies the sublog parameter being transferred as the Number of Blocks Received sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The NUMBER OF BLOCK RECEIVED field contains the number of blocks received from the service deliver subsystem for the device server of the addressed logical unit as a result of write commands (see 1.0.1.1).

Table 6 shows the format of Number of Blocks Sent sublog parameter.

Table 7 — Number of Blocks Sent sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0004h)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	NUMBER OF BLOCKS SENT							(LSB)

The PARAMETER CODE field set to 0004h identifies the sublog parameter being transferred as the Number of Blocks Sent sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The NUMBER OF BLOCK SENT field contains the number of blocks delivered to the service deliver subsystem by the device server of the addressed logical unit as a result of read commands (see 1.0.1.1).

Table 8 shows the format of Read Command Processing Time sublog parameter.

Table 8 — Read Command Processing Time sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0005h)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	READ COMMAND PROCESSING TIME							(LSB)

The PARAMETER CODE field set to 0005h identifies the sublog parameter being transferred as the Read Command Processing Time sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The READ COMMAND PROCESSING TIME field contains the cumulative amount of time spent processing read commands addressed to logical unit (see 1.0.1.1).

Table 9 shows the format of Write Command Processing Time sublog parameter.

Table 9 — Write Command Processing Time sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0006h)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	WRITE COMMAND PROCESSING TIME							(LSB)

The PARAMETER CODE field set to 0006h identifies the sublog parameter being transferred as the Write Command Processing Time sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The WRITE COMMAND PROCESSING TIME field contains the cumulative amount of time spent processing write commands addressed to logical unit (see 1.0.1.1).

Table 10 shows the format of Weighted Number of Read Commands Plus Write Commands sublog parameter.

Table 10 — Weighted Number of Read Commands Plus Write Commands sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____ PARAMETER CODE (0007h) _____ (LSB)							
1								
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____ WEIGHTED NUMBER OF READ COMMANDS PLUS WRITE COMMANDS _____ (LSB)							
7								

The PARAMETER CODE field set to 0007h identifies the sublog parameter being transferred as the Weighted Number of Read Commands Plus Write Commands sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

If task priority is supported (see SAM-4), then the WEIGHTED NUMBER OF READ COMMANDS PLUS WRITE COMMANDS field contains the cumulative weighted number of read commands and write commands addressed to logical unit (see 1.0.1.1).

If task priority is not supported, then the WEIGHTED NUMBER OF READ COMMANDS PLUS WRITE COMMANDS field shall be set to zero.

Table 11 shows the format of Weighted Read Command Processing Plus Write Command Processing sublog parameter.

Table 11 — Weighted Number of Read Command Processing Plus Write Command Processing sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____ PARAMETER CODE (0008h) _____ (LSB)							
1								
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____ WEIGHTED NUMBER OF READ COMMAND PROCESSING PLUS WRITE COMMAND _____ (LSB)							
7	PROCESSING							

The PARAMETER CODE field set to 0008h identifies the sublog parameter being transferred as the Weighted Number of Read Command Processing Plus Write Command Processing sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

If task priority is supported (see SAM-4), then the WEIGHTED NUMBER OF READ COMMAND PROCESSING PLUS WRITE COMMAND PROCESSING field contains the cumulative weighted amount of time spent processing read commands and write commands addressed to logical unit (see 1.0.1.1).

If task priority is not supported, then the WEIGHTED NUMBER OF READ COMMAND PROCESSING PLUS WRITE COMMAND PROCESSING field shall be set to zero.

Table 12 shows the format of the Idle Time sublog parameter.

Table 12 — Idle Time sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0009h)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	IDLE TIME							(LSB)

The PARAMETER CODE field set to 0009h identifies the sublog parameter being transferred as the Idle Time sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The IDLE TIME field contains the cumulative amount of time spent when there are no tasks in the task set and there are no tasks being processed by a logical unit.

Table 13 shows the format of the Time Interval sublog parameter.

Table 13 — Time Interval sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (000Bh)							(LSB)
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (08h)							
4	(MSB) _____							
11	TIME INTERVAL							(LSB)

The PARAMETER CODE field set to 000Bh identifies the sublog parameter being transferred as the Time Interval sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 14.

The PARAMETER LENGTH field shall contain 08h.

The TIME INTERVAL field contains the time interval in nanoseconds used in the Read Command Processing Time sublog page, the Write Command Processing Time sublog page, and the Weight Read Command Processing Plus Write Command Processing sublog page.

1.0.1.3 Group Statistics and Performance subpage log

Table 14 specifies the General Statistics and Performance subpage log parameters.

Table 14 — Group Statistics and Performance log subpages

Bit Byte	7	6	5	4	3	2	1	0
0	PAGE CODE (xxh)							
1	SUBPAGE CODE							
2	(MSB)	PAGE LENGTH (30h)						(LSB)
3								
Group Statistics and Performance sublog parameters								
4	Group n Number of Read Commands sublog parameter							
11								
12	Group n Number of Write Commands sublog parameter							
19								
20	Group n Number of Blocks Received sublog parameter							
27								
28	Group n Number of Blocks Sent sublog parameter							
35								
36	Group n Read Command Processing Time sublog parameter							
43								
44	Group n Write Command Processing Time sublog parameter							
51								

The PAGE CODE and PAGE LENGTH fields are described in 7.2.1.

The SUBPAGE CODE field is as specified in table 2.

Table 15 shows the format of Group n Number of Read Commands sublog parameter.

Table 15 — Group n Number of Read Commands sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0001h)						(LSB)
1								
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB)	GROUP N NUMBER OF READ COMMANDS						(LSB)
7								

The PARAMETER CODE field set to 0001h identifies the sublog parameter being transferred as the Group n Number of Read Commands sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N NUMBER OF READ COMMANDS field contains the number of read commands (see 1.0.1.1) received by the addressed logical unit.....

Table 16 shows the format of Group n Number of Write Commands sublog parameter.

Table 16 — Group n Number of Write Commands sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0002h)							(LSB)
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	GROUP N NUMBER OF WRITE COMMANDS							(LSB)

The PARAMETER CODE field set to 0002h identifies the sublog parameter being transferred as the Group n Number of Write Commands sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N NUMBER OF WRITE COMMANDS field contains the number of write commands (see 1.0.1.1) received by the addressed logical unit.

Table 6 shows the format of Group n Number of Blocks Received sublog parameter.

Table 17 — Group n Number of Blocks Received sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0003h)							(LSB)
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	GROUP N NUMBER OF BLOCKS RECEIVED							(LSB)

The PARAMETER CODE field set to 0003h identifies the sublog parameter being transferred as the Group n Number of Blocks Received sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N NUMBER OF BLOCK RECEIVED field contains the number of blocks received from the service deliver subsystem for the device server of the addressed logical unit as a result of write commands (see 1.0.1.1).

Table 6 shows the format of Group n Number of Blocks Sent sublog parameter.

Table 18 — Group n Number of Blocks Sent sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0004h) _____ (LSB)							
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	GROUP N NUMBER OF BLOCKS SENT _____ (LSB)							

The PARAMETER CODE field set to 0004h identifies the sublog parameter being transferred as the Group n Number of Blocks Sent sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N NUMBER OF BLOCK SENT field contains the number of blocks delivered to the service deliver subsystem by the device server of the addressed logical unit as a result of read commands (see 1.0.1.1).

Table 8 shows the format of Group n Read Command Processing Time sublog parameter.

Table 19 — Group n Read Command Processing Time sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0005h) _____ (LSB)							
2	DU	DS	TSD	ETC	TMC	LBIN	LP	
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	GROUP N READ COMMAND PROCESSING TIME _____ (LSB)							

The PARAMETER CODE field set to 0005h identifies the sublog parameter being transferred as the Group n Read Command Processing Time sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N READ COMMAND PROCESSING TIME field contains the cumulative amount of time spent processing read commands addressed to logical unit (see 1.0.1.1).

Table 9 shows the format of Group n Write Command Processing Time sublog parameter.

Table 20 — Group n Write Command Processing Time sublog parameter format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
1	PARAMETER CODE (0006h) _____ (LSB)							
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (04h)							
4	(MSB) _____							
7	GROUP N WRITE COMMAND PROCESSING TIME _____ (LSB)							

The PARAMETER CODE field set to 0006h identifies the sublog parameter being transferred as the Group n Write Command Processing Time sublog parameter.

The values of the log parameter control bits for self test results log parameters is specified in table 4.

The PARAMETER LENGTH field shall contain 04h.

The GROUP N WRITE COMMAND PROCESSING TIME field contains the cumulative amount of time spent processing write commands addressed to logical unit (see 1.0.1.1).