

Logical Unit Self-tests

To: T10 Technical committee
From: Mark Evans
Quantum Corporation
500 McCarthy Boulevard
Milpitas, CA USA 95035
Phone: 408-894-4019
Fax: 408-952-3620
Email: mark.evans@quantum.com
Date: 08 July 1999

1 Introduction

As the complexity of today's systems increases, a method is needed to quickly identify the root cause of system problems at the end-user's site. Self-test routines implemented in a logical unit (e.g., a hard disk drive) would provide a mechanism to quickly determine whether the logical unit is faulty or indicate that some other component is causing the system to fail. These self-test routines could be performed by a logical unit without removing it from its installation. This proposal recommends changes to SPC-x to provide these features.

In revision 2 of this proposal, I've tried to include all comments, issues and responses that I've received to date. In addition, at the request of the editor of SPC-2, I've tried to make this proposal generic so that it could apply to any type of device that could be a logical unit. I would appreciate input on issues that are raised in this revision. Please feel free to continue supplying comments to me or the SCSI reflector as appropriate.

2 Reference documents

Elements from the draft standard SCSI Primary Commands – 2 (SPC-2) Revision 9, 15 March 1999 are referenced in this proposal. For more detail, see the referenced clauses. This draft is available at <ftp://ftp.symbios.com/pub/standards/io/t10/drafts/spc2-r09.pdf>.

3 Overview model

3.1 The short and extended self-tests

Two types of self-tests are defined in this proposal. One is a short self-test. One is an extended self-test. The primary goal of the short self-test is to quickly identify if the logical unit determines that it is faulty. The primary goal of the extended self-test routine is to simplify factory testing during integration by having logical units perform more comprehensive testing without application client intervention. A second goal of the extended self-test is to provide a more comprehensive test to validate the results of the short self-test if its results are judged by the application client to be inconclusive.

The criteria for the short self-test are that the self-test has one or more segments and completes in two minutes or less. The criteria for the extended self-test are that it has one or more segments and that the completion time is vendor-specific. Any tests performed in the segments are vendor-specific.

Some possible segments might be:

- 1) An electrical segment wherein the logical unit tests its own electronics. Though the particular tests in this segment are vendor specific, this segment might include such tests as a buffer RAM test, a read/write circuitry test, and/or a test of the read/write head elements.
- 2) A seek/servo segment wherein a device tests its capability to find and servo on data tracks.
- 3) A read/verify scan segment wherein a device performs read scanning of some or all of the medium surface.

The tests performed in the segments may be the same for the short and extended self-tests. The time required by a logical unit to perform its extended self-test is reported in the EXTENDED SELF-TEST ROUTINE COMPLETION TIME field in the Control mode page as defined in this proposal.

3.2 Self-test modes

There are two modes for each type of self-test (short and extended): a background mode and a foreground mode. In the background mode executing the self-test routine has lower priority than servicing any other commands. While performing a self-test in the background mode the device server shall suspend the self-test and respond to all commands within two seconds. There should be negligible performance degradation for other commands while the self-test routine is running in the background mode. **[EDITOR'S NOTE: The previous sentence may sound like marketing literature, and I'm willing to delete it. However, I was trying to provide as many tips to implementers as possible.]**

While performing a self-test in the foreground mode the device server shall respond to all commands except INQUIRY, LOG SENSE, REPORT LUNS, and REQUEST SENSE with CHECK CONDITION status with a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY, OPERATION IN PROGRESS. It is recommended that, before an application client initiates a self-test in the foreground mode that it should reserve the logical unit.

[EDITOR'S NOTE: There was some discussion about including START/STOP UNIT (with the START bit set to one and the POWER CONDITIONS field containing a value of zero), but this has been removed. This is fine by me. Is there other input here?]

REQUEST SENSE may be used by the application client at any time during execution of a self-test to poll the logical unit's progress. The logical unit shall return the percent of self-test completion in the PROGRESS INDICATION field in the sense data when the Self-test progress indication bit (STPI) in the REQUEST SENSE CDB is set to one.

After the PROGRESS INDICATION field returned in response to a REQUEST SENSE command indicates that the self-test is complete, the application client should use LOG SENSE to determine the self-test results for a self-test that was performed in the background mode.

The following table describes when a logical unit returns status after receipt a self-test command, how an application client may abort a self-test, how an logical unit handles new commands that are received while a self-test is in progress, and how a logical unit reports a self-test failure.

Mode	When Status is Returned	How to abort the self-test	Action upon receipt of new commands	Self-test failure reporting
Fore-ground	After the self-test is complete	Abort Task message	CHECK CONDITION / NOT READY/ LOGICAL UNIT NOT READY, DIAGNOSTICS IN PROGRESS	CHECK CONDITION / HARDWARE ERROR/ LOGICAL UNIT FAILED SELF-TEST
Back-ground	After the command is received	SEND DIAGNOSTICS w/ FUNCTION CODE field = 100h	Process the command	Check self-test results log page after the PROGRESS INDICATION field returned from REQUEST SENSE indicates self-test complete

[EDITOR'S NOTE: There has been much discussion about the content of the table above. Originally, it was thought to include an IMMED bit, but this was eliminated, so that we wouldn't have to deal with the deferred errors.]

3.3 Test sequence

The self-tests and modes are invoked by the application client by issuing a SEND DIAGNOSTICS command to the logical unit with the FUNCTION CODE field set to the appropriate value. If the value in the FUNCTION CODE field specified a self-test in foreground mode, the device server shall report status, set the function code value into the FUNCTION CODE field and set the SELF-TEST RESULTS VALUE field to be Fh in its Self-test results log page, store the log page (to non-volatile memory, if possible), and begin the first self-test segment. If the value in the FUNCTION CODE field specified a self-test in background mode, the device server shall set the function code value into the FUNCTION CODE field and set the SELF-TEST RESULTS VALUE field to be Fh in its Self-test results log page, store the log page (to non-volatile memory, if possible), and begin the first self-test segment.

If three self-test segments are implemented, the device server shall change the value in the SELF-TEST RESULTS VALUE field to be Eh after completing the first two self-test segments of the self-test, update the log page (store the page in non-volatile memory, if possible), begin the third self-test segment. Only the SELF-TEST RESULTS VALUE field shall be changed as a result of this update. Updating the log allows the application client to read the Self-test Results log page to determine that the third segment is in progress.

The application client may query the device server for progress by using the REQUEST SENSE command with the STPI bit set to one (see 5.1). The device server shall then report its progress through the test in the PROGRESS INDICATION field.

[EDITOR'S NOTE: Use of the RECEIVE DIAGNOSTIC RESULTS command to poll for progress was considered and was rejected in favor of REQUEST SENSE for at least three reasons: 1) it was thought that one RECEIVE DIAGNOSTIC RESULTS needed to be paired with one SEND DIAGNOSTICS command so that an application client would be unable to poll multiple times

during a self-test; 2) the description of the SEND DIAGNOSTICS command in SPC-2 says that this command is usually followed by a RECEIVE DIAGNOSTIC RESULTS command *except following a self-test*; and, 3) it was thought that most devices would handle a REQUEST SENSE command more easily than a RECEIVE DIAGNOSTIC RESULTS command while performing a self-test in the foreground mode.]

After a self-test is completed or aborted the device server shall update the SELF-TEST RESULTS VALUE field in its Self-test results log page in non-volatile memory. The application client may use LOG SENSE to read the results from up to the last twenty self-tests performed by the device server.

If a self-test is in progress in the foreground mode and an error occurs, the device server shall update the fields in its Self-test results log page and report CHECK CONDITION / HARDWARE ERROR/ LOGICAL UNIT FAILED SELF-TEST.

4 Mode page entry

The following describes the EXTENDED SELF-TEST ROUTINE COMPLETION TIME field:

[EDITOR'S NOTE: A mode page needs to be identified where this could be placed. It is recommended that bytes 10 and 11 (currently "reserved") in the Control mode page be used for this (see 8.3.4 in SPC-2). The following proposal assumes that this recommendation is accepted.]

Table 1 – Control mode page

Bit	7	6	5	4	3	2	1	0	
Byte									
.....								
10	(MSB)	EXTENDED SELF-TEST ROUTINE COMPLETION TIME							
11								(LSB)	

The EXTENDED SELF-TEST ROUTINE COMPLETION TIME is an advisory parameter that the application client may use to determine the time in seconds that the device server requires to complete the extended self-test routine when the device server is not interrupted by an application client and no errors occur during execution of the routine. The application client should expect this time to increase significantly if other commands are sent to the target or if errors occur during execution of the routine.

5 Modifications to the REQUEST SENSE command

The following are two proposed modifications to the REQUEST SENSE command in SPC-2.

5.1 Self-test progress indication bit

The following describes the inclusion of a new bit in the REQUEST SENSE command that specifies that a self-test progress indication be returned to the application client as a response to this command.

[EDITOR'S NOTE: The changes in this clause should be included in 7.20 REQUEST SENSE command in SPC-2]

Table 2 – REQUEST SENSE command

Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (03h)							
1	RESERVED							STPI
2	RESERVED							
3	RESERVED							
4	ALLOCATION LENGTH							
5	CONTROL							

If the device server supports the self-test function, the device server shall return the percent of self-test completion in the PROGRESS INDICATION field in the sense data when the Self-test progress indication bit (STPI) is set to one. When STPI is cleared to zero the device server may return the PROGRESS INDICATION field.

5.2 Progress indication field

[EDITOR’S NOTE: The following is the text and note currently in SPC-2 in clause 7.20.1 following Table 66 — Progress indication bytes].

The PROGRESS INDICATION field is a percent complete indication in which the returned value is the numerator that has 65536 (10000h) as its denominator. The progress indication shall be based upon the total operation.

NOTE 28 It is intended that the progress indication be time related. However, since for example format time varies with the number of defects encountered, etc., it is reasonable for the device server to assign values to various steps within the process. The granularity of these steps should be small enough to provide reasonable assurances to the application client that progress is being made.

[EDITOR’S NOTE: The should be included in 7.20.1 Sense-key specific following the above text].

If the device server is responding to a REQUEST SENSE command where the STPI bit is set to one and the most recent self-test is complete, then the PROGRESS INDICATION field shall contain 0000h.

6 Modifications to the SEND DIAGNOSTICS command

The last line in the first paragraph of this clause in SPC needs to be changed to, “Except when the SELFTEST bit is one or the FUNCTION CODE field is not 000b, this command is usually followed by a RECEIVE DIAGNOSTIC RESULTS command (see 7.15).”

For this proposal reserved bits 7, 6 and 5 in byte 1 of the SEND DIAGNOSTICS command are changed to a FUNCTION CODE field, and reserved bit 3 in byte 1 is changed to an IMMED bit. The following table and text describes the proposed changes for this command.

Table 3 – SEND DIAGNOSTICS command

Bit	7	6	5	4	3	2	1	0
Byte								
0	OPERATION CODE (1Dh)							
1	FUNCTION CODE			PF	RESERVED	SELFTEST	DEVOFFL	UNITOFFL
2	RESERVED							
3	(MSB)	PARAMETER LIST LENGTH						(LSB)
4								
5	CONTROL							

When the SELFTEST bit is set to one the FUNCTION CODE field shall be 000b. When the SELFTEST bit is cleared to zero the content of FUNCTION CODE field is specified in the following table.

Table 4 – Function code field values

Value	Description
000b	Value to be used when bit 2 of this byte is set to one or if the SEND DIAGNOSTIC command is not invoking one of the other self-test function codes (e.g., Translate Address page (see SBC)).
001b	The device server shall start its short self-test routine in the background mode (see note 2).
010b	The device server shall start its extended self-test routine in the background mode (see note 2).
011b	Reserved
100b	Abort the current short or extended self-test routine running in background mode. This value is only valid if a previous SEND DIAGNOSTIC command requested a Start short or extended self-test routine in background mode has not completed. If these conditions are not true, then the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
101b	The device server shall start its short self-test routine in the foreground mode (see note 3).
110b	The device server shall start its extended self-test routine in the foreground mode (see note 3).
111b	Reserved
NOTES –	
1) When a self-test with a Function Code other than 000b is completed or aborted the device server shall log the self-test results in the Self-test results log page.	
2) In the background mode the self-test routine shall have lower priority than servicing any other commands. The device server may pause execution of its self-test routine to service any other command. While performing a self-test in the background mode the device server shall respond to any command within two seconds.	
3) While performing a self-test in the foreground mode the device server shall respond to all commands except INQUIRY, LOG SENSE, REPORT LUNS, and REQUEST SENSE with CHECK CONDITION status. The sense key shall be set to NOT READY (2h) with the additional sense code set to LOGICAL UNIT NOT READY, DIAGNOSTICS IN PROGRESS. [EDITOR'S NOTE: a new ASC/ASCQ needs to be assigned for LOGICAL UNIT NOT READY, DIAGNOSTICS IN PROGRESS. It is recommended that this code be 04h,09h.]	

Device servers supporting FUNCTION CODE field values other than 000b for the SEND DIAGNOSTIC command, shall support the EXTENDED SELF-TEST ROUTINE COMPLETION TIME field in the Control mode page.

7 Self-test results log page

This page provides the results from the twenty most recent self-tests. Results from the most recent self-test or the self-test currently in progress shall be reported in the first self-test log structure parameter; results from the second most recent self-test shall be reported in the second self-test log structure parameter; etc. If fewer than twenty self-tests have occurred, the unused self-test log structure parameter entries shall be zero filled.

The following table describes the Self-test results log page format that shall be returned by the devices server upon request by the application client.

Table 5 – Self-test results log page format

Byte	Bit	7	6	5	4	3	2	1	0
0		PAGE CODE (xxh)							
1		RESERVED							
2	(MSB)	PAGE LENGTH (0xxxh)							
3		(LSB)							
SELF-TEST RESULTS LOG PARAMETERS									
4		FIRST SELF-TEST RESULTS LOG PARAMETER							
.....									
19									
.....									
307		TWENTIETH SELF-TEST RESULTS LOG PARAMETER							
.....									
323		(LSB)							

The PAGE CODE field shall be xxh.

[EDITOR'S NOTE: a number needs to be assigned to this log page.]

The PAGE LENGTH field specifying the length in bytes of the parameter list that shall be transferred from the device server to the application client shall be 0xxxh.

The following table describes the nth SELF-TEST RESULTS LOG PARAMETER field.

Table 6 – nth self-test results log parameter

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE							
1		(LSB)							
2		DU (= 0)	DS (= 0)	TSD (= 0)	ETC (= 0)	TMC (= 00b)	LBIN (= 1)	LP (= 1)	
3		PARAMETER LENGTH (08h)							
4	RSVD	FUNCTION CODE			SELF-TEST RESULTS VALUE				
5		RESERVED							
6	(MSB)	TIMESTAMP							
7		(LSB)							
8	(MSB)	LBA OF FIRST FAILURE							
.....									
11		(LSB)							

The FUNCTION CODE field contains the content of the FUNCTION CODE field in the SEND DIAGNOSTICS command that initiated this self-test. See Table 4 for a description of the FUNCTION CODE field values.

[EDITOR'S NOTE: The FUNCTION CODE field was reduced to three bits to be the same size as the FUNCTION CODE field in the SEND DIAGNOSTICS command.]

The following table defines the content of the SELF-TEST RESULTS VALUE field.

Table 7 – Self-test results values

Value	Description
0h	The self-test routine completed without error.
1h	The self-test routine was aborted by the application client.
2h	The self-test routine was interrupted by a reset.
3h	An unknown error occurred while the device server was executing the self-test routine and the device server was unable to complete the self-test routine.
4h	The self-test completed having a test element that failed, and the test element that failed is not known.
5h	The first segment of the self-test failed.
6h	The second segment of the self-test failed.
7h	The third segment of the self-test failed.
8h-Dh	Reserved
Eh	The third segment of the self-test is in progress
Fh	The first or second segment of the self-test is in progress

If the logical unit implements logical blocks, the content of the LBA OF FIRST FAILURE field is the first logical block address where a self-test error occurred. This value shall not change (e.g., as the result of block reassignment). The content of the LBA OF FIRST FAILURE field shall be FFFFFFFh if no errors occurred during the self-test or if the error that occurred is not related to a logical block address.

The TIMESTAMP field contains the total accumulated power-on hours of the device server at the time the self-test operation was completed. If the test is still in progress, the content of the TIMESTAMP field shall be zero. If the power-on hours for the device server at the time the self-test operation was completed is greater than FFFFh then the content of the TIMESTAMP field shall be FFFFh.