

To: T10 Committee
From: Gerry Houlder, Seagate Technology <gerry.houlder@seagate.com>
Subj: Background Medium Scan
Date: June 30, 2004

This document proposes a standard method to control and retrieve status from background medium scan operations. Several drive vendors (including Seagate) have proprietary methods for controlling these operations today, but customers have asked that a standard method be defined so their processes can be simpler.

Many systems implement background read scanning in their operating systems. Having the drive do the read scanning instead has the advantage of reducing system overhead and removing an unproductive demand on interface bus bandwidth. During idle periods, the drive reads each sector on its medium to verify the data. If correctable errors are detected, the drive may rewrite and/or relocate the LBA to another sector depending on what has been allowed by the customer. Errors are logged to a list; if allowed by the customer, a future write to that LBA may result in the LBA being relocated before performing the write.

There is a related option called Pre-scan. This is a feature that starts a background medium scan right after power on. If the customer does a write operation to an area of the disk that has not been scanned yet, the drive turns the write operation into a write-and-verify operation. This has the advantage of checking (right after power on) that all data can be read back correctly. A disadvantage is that the drive's performance is greatly reduced (because of the extra steps to do the verify operation) until the first medium scan is completed. Some customers choose to disable this option for that reason.

This proposal includes a mode page that is used to control the scan, pre-scan, and related features. It also includes a log page that stores a list of LBAs that have or did have uncorrectable read errors. This provides a list of LBAs that the drive knows are at risk and the system may not know about yet.

1.0 Background medium scan model

[not written yet]

2.0 BMS Control page

The background medium scan (BMS) control page (see table 1) is used to control the background medium scan feature.

Table 1 – BMS Control Page

Bit	7	6	5	4	3	2	1	0	
0	PAGE CODE (??h)								
1	PAGE LENGTH (0Ah)								
2								EN_PS	
3								EN_BMS	
4	(MSB)	BMS INTERVAL TIME							
5								(LSB)	
6	(MSB)	PRESCAN TIMEOUT VALUE							
7								(LSB)	
8	RESERVED								
9	RESERVED								
10	RESERVED								
11	RESERVED								

An enable pre-scan (EN_PS) bit set to zero specifies that pre-scan is disabled. If a pre-scan operation is in progress when EN_PS is changed from a one to a zero then pre-scan is halted. An EN_PS bit set to one specifies that a pre-scan operation is started after the next power on cycle. Once this pre-scan has completed, another pre-scan shall not occur unless the EN_PS bit is set to zero, then set to one, and another power on cycle occurs.

An enable background medium scan (EN_BMS) bit set to zero specifies that background medium scan is disabled. An EN_BMS bit set to one then background medium scan operations are allowed. If a background medium scan is in progress when the EN_BMS bit is changed from one to zero, then the medium scan is suspended; when the EN_BMS bit is subsequently set to one, the medium scan shall resume from the suspended location.

The BMS interval time field specifies the minimum time, in hours, between the end of one medium scan operation and the start of the next medium scan operation. The minimum value for this time is two hours.

The pre-scan timeout field specifies the maximum time, in hours, for a pre-scan operation to complete. If the pre-scan operation doesn't complete within the specified time then it is halted. The minimum value for this time is two hours.

3.0 Medium Scan log page

The Medium Scan log page (see table 2) contains the BMS Status parameter and zero or more Medium Error log entries. The BMS Status parameter provides information about pre-scan and background medium scan operations. Each medium error log entry corresponds to a site where an error was detected. If the Medium Scan log area is filled up, a new medium error log entry will overwrite the oldest entry. When a LOG SELECT command with PCR bit set to one is processed the Medium Error log entries are deleted, but the values in the BMS Status Parameter are not affected.

Table 2 – Medium Scan Log Page

Byte	Bit	7	6	5	4	3	2	1	0
0		PAGE CODE (??h)							
1		RESERVED							
2	(MSB)	PAGE LENGTH (N-3)							
3		(LSB)							
4		BMS STATUS PARAMETER							
23		BMS STATUS PARAMETER							
24		FIRST MEDIUM ERROR LOG PARAMETER (LEAST RECENT)							
		ADDITIONAL MEDIUM ERROR LOG PARAMETERS							
N		LAST MEDIUM ERROR LOG PARAMETER (MOST RECENT)							

3.1 BMS Status parameter

The background medium scan (BMS) status parameter (see table 3) contains status information about the background medium scan feature.

Table 3 – BMS Status parameter format

Bit	7	6	5	4	3	2	1	0
Byte								
0	(MSB) PARAMETER CODE (0000h)							
1	(LSB)							
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (10h)							
4	(MSB)							
5	TIME STAMP							
6								
7	(LSB)							
8	RESERVED							
9	BMS STATUS							
10	(MSB) NUMBER OF SCANS PERFORMED							
11	(LSB)							
12	(MSB)							
13	CURRENT LBA FOR SCAN							
14								
15								
16								
17								
18								
19	(LSB)							

The TIME STAMP indicates the drive power on minutes at the time the log page is requested.

Table 4 specifies the meaning of the BMS STATUS field.

Table 4 – BMS Status field

Value	Meaning of definition parameter
0h	No scans active.
1h	Background medium scan is active.
2h	Pre-scan is active.
3h	Fatal error, scan halted.
4h	Scan suspended by the drive.
5h	Scan suspended by initiator request.
6h - FFh	Reserved

The NUMBER OF SCANS PERFORMED field indicates the number of background scans that have been performed since the drive was originally shipped by the drive vendor.

The CURRENT LBA FOR SCAN field indicates the next LBA that will be scanned when a pre-scan or background medium scan is active. When no scans are active, the value is set to zero.

Table 5 – Medium Scan parameter format

Byte	7	6	5	4	3	2	1	0
0	(MSB) PARAMETER CODE (????h through ????) (LSB)							
1								
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (14h)							
4	(MSB) TIME STAMP							
5								
6								
7	(LSB)							
8	REASSIGN STATUS				SENSE KEY			
9	ADDITIONAL SENSE CODE							
10	ADDITIONAL SENSE CODE QUALIFIER							
11	HEAD NUMBER				(MSB)	CYLINDER NUMBER		
14	CYLINDER NUMBER							
15	(LSB)							
14	(MSB) SECTOR NUMBER							
15	(LSB)							
16	(MSB) CURRENT LBA FOR SCAN							
17								
18								
19								
20								
21								
22								
23	(LSB)							

The TIME STAMP indicates the drive power on minutes at the time the error was detected.

Table 6 specifies the meaning of the REASSIGN STATUS field.

Table 6 – Reassign Status field

Value	Meaning of definition parameter
0h	No reassignment needed.
1h	Reassignment pending receipt of Reassign command or write command (if auto write reallocation is allowed) from the initiator.
2h	LBA successfully reassigned by drive.
3h	reserved
4h	Reassignment failed.
5h	LBA recovered via re-write only.
6h - Fh	Reserved

The Sense Key, Additional Sense Code, and Additional Sense Code Qualifier fields are the same as the sense bytes (see Request Sense command). They indicate the sense bytes that would be reported for this site if the medium error were detected by a normal read or write. These sense

bytes may not have been reported because the error was detected during background scan or reporting of that type of error was disabled by the initiator.

The head number field indicates the head on which the medium error was detected.

The cylinder number field indicates the cylinder number associated with the medium error.

The sector number field indicates the sector number associated with the medium error.

The current LBA for scan field indicates the next LBA to be accessed during pre-scan or background medium scan activity. If pre-scan and background scan are not active, this value is set to zero.