

To: X3T9.2 Membership
From: Gregory G. Floryance
IBM Corporation
Low End Storage Products
System Products Division
Hwy 52 & 37th Street NW
Rochester, Mn 55901
(507) 253-1869

Date: 10/22/87

Subject: Format Progress Reporting (Document 87-180 Rev 1)

PROBLEM STATEMENT.....

Typically, the Format command requires a great deal of time to execute. During command execution, it is currently not possible for an Initiator to determine either command progress or anticipated time to completion. For the system operator, the lack of assurance that such information provides can lead to anxiety and an unfortunate response such as a power-down cycle.

Historically, PC-DOS has provided operator feedback by indicating the cylinder and head number being formatted. In the future, larger DASD devices will consume even greater time to format which will result in a greater need for progress report capability. To achieve a user friendly environment, SCSI should standardize a progress reporting method to both the issuing initiator as well as other interested initiators.

One mechanism for implementing progress reporting which was considered and rejected was to send a messages on a per cylinder crossing boundary to the initiator. The initiator would then increment a counter and display the value. One problem was that messages are used to control the physical layer of SCSI and are not intended to be surfaced to the host operating system. A second problem was that it requires the SCSI bus to be dedicated (no disconnect) so that the target does not have to compete with other devices for the bus each time a cylinder seek was required. A third problem was in the assumption that all devices were physically similar in the amount of time spent on a cylinder. A single surface device could cross 50 cylinders a second while a multiple track device may cross 1 cylinder per second. A final problem was this technique offered no distinction for reporting certification progress.

A far better proposal which meets all the command requirements within the architectural constraints of SCSI is described in the following proposal.

PROPOSAL SUMMARY.....

This proposal defines an Immed bit in the Format CDB which indicates to the target that following the verification of the CDB and parameter data, GOOD status and COMMAND COMPLETE should be returned prior to the start of the format operation. This is similar to the Immed bit of the START/STOP UNIT command. This proposal also includes modifications to the Request Sense command so that the format progress information can be returned to all initiators.

DETAILED PROPOSAL.....

Additional sense code FORMAT IN PROGRESS defined for NOT READY sense key.

Format command modifications to define one of the Reserved bits in either the FORMAT CDB or in the DEFECT LIST DATA to be the Immed bit. Add the following to section 8.1.1 to describe the Immed bit function:

An immediate (Immed) bit value of one indicates that the status shall be returned as soon as the operation is initiated. Following the successful verification of the content of the CDB and the Defect List data, the target shall return GOOD status followed by COMMAND COMPLETE. During the format operation, the target shall return a sense key of NOT READY and an additional sense code of FORMAT IN PROGRESS. For those commands which received a prior CHECK CONDITION as a result of the format operation, optional format in progress information may also be available in the sense key specific bytes as described in section 7.1.8.1.

An Immed bit of zero indicates that status shall be returned after the operation is completed.

Request Sense Modifications in section 7.1.8.1 to add the following.

: If the sense key is NOT READY and the additional sense code is FORMAT IN PROGRESS and the SKSV bit is one, then the sense key specific bytes shall be defined as shown in Table 7-23b. Bytes 16-17 shall contain a percent complete indication in which the returned value is the numerator that has 256 as it's denominator. The percent complete shall be based upon the total format operation including any certification or initialization operations.

Note: It is intended that the progress indication be time related. However, since format time varies with the number of defects encountered, etc., it is reasonable for the target to assign values to various steps within the process. The granularity of these steps should be small enough to provide reasonable assurances to the initiator that progress is being made.

Note: In order to receive accurate data, the returned sense data should reflect the most recent progress of the format operation and not that which existed at the time that a previous command returned CHECK CONDITION status.

	7	6	5	4	3	2	1	0	
15	SKSV		Reserved						
16	(MSB)		Format Progress Indication						
17	(LSB)								

Figure 1. Table 7-23b: Format Progress Indication

103

REQUESTED ENHANCEMENTS.....

After having submitted the original proposal, a number of comments and modifications were requested by various individuals. These include the following:

- Define bits to distinguish between Filling, Certifying, Formatting
- Return CURRENT CYLINDER along with TOTAL CYLINDERS
- Return CURRENT PASS along with TOTAL PASSES
- Return CURRENT HEAD number

In the discussions which arose following these requests, it became obvious that during format, the technique (process) used was extremely device dependent. Some devices format, certify and initialize on a per cylinder basis. Other devices completed one of these operations for all cylinders before starting the next operation. Other devices did a combination of both. The most extreme example mentioned was a mapping of multiple devices with differing physical characteristics (cylinders) into a single LUN.

WHY ENHANCEMENTS SHOULD NOT BE INCORPORATED.....

As a result of the device uniqueness's, any physical information which could be returned would only lead to confusion unless the format process itself was standardized. How can an operator estimate the amount of time to format if the process for formatting is not known? If a consistent interpretation by the initiator is not possible, then the fields become useless and should not be defined at all. As a result of these discussions, it becomes evident that the only practical common denominator is that of reporting a percent complete.

A second argument against the incorporation of the last three enhancements arises from the fact that it is currently impossible to accommodate the remaining requests given the current definition of the Request Sense command data fields. The INFORMATION BYTES (3-6) and the COMMAND SPECIFIC INFORMATION BYTES (8-11) are both defined as command specific. Since it is the goal of this proposal to report format progress (in both single and multi-initiator systems) without being coupled to an outstanding command, neither of these fields can be used unless the definition of these fields is changed.

HOWEVER, IF ENHANCEMENTS ARE ADOPTED.....

The first item could be incorporated into this proposal by using the remainder of byte 15 as shown in Figure 2 on page 4. The following definitions could then apply:

The Fill bit indicates that the data fields of the medium are currently being initialized with the either initiator specified initialization data pattern or with a data pattern defined by the target.

The Ctfy bit indicates that the medium is currently being certified.

The Glst bit indicates that the G list defects are currently being installed on the medium.

The Fmt bit indicates that medium is currently being formatted.

Note: More than one bit of byte 15 may be active at one time.

The last three items could be incorporated by changing the definition of the Request Sense Information Bytes and Command Specific Information Bytes so that they are not restricted to a command but rather to a sense key/code. The following field definitions could then apply:

The INFORMATION BYTES 3-6 are optional. Their content is considered valid if the VALID bit (byte 0, bit 1) is one. For a VALID bit of zero, bytes 3-6 shall be zero. Bytes 3-4 shall contain the current cylinder number currently being formatted. Bytes 5-6 shall contain the Total Cylinders.

The COMMAND SPECIFIC INFORMATION BYTES 8-11 are optional. Bytes 8-9 shall contain the current pass. Bytes 10-11 shall contain the total number of passes.

0	Valid	Error Code			
3	(MSB)	Current Cylinder			
4	(LSB)				
5	(MSB)	Total Cylinders			
6	(LSB)				
8	(MSB)	Current Pass			
9	(LSB)				
10	(MSB)	Total Pass			
11	(LSB)				
15	SKSV	Reserved		Fill	Ctfy Glst Fmt

104