

TEXAS INSTRUMENTS



MEMORANDUM -- 01 DEC 1989

D. W. (Bill) Spence

Member, Group Technical Staff
Computer Systems Division
Data Systems Group

TO: John Lohmeyer, Chairman X3T9.2

FROM: Bill Spence, Texas Instruments

12501 Research Boulevard • P.O. Box 2909 • MS 2091 • Austin, Texas 78769
512 250-6627 • Fax 512 250-7479 • TWX 910 867-4702 • Telex 73324

SUBJECT: Reflections on Writing to EOT (for Mag Tape people)

In section 9.2.14 of SCSI-2, rules are given for the behavior of a tape drive if it approaches the end of a partition while executing a write command. The rules are workable and acceptable for the standard. They rather smother one related concept which was aired in debate in back in 1987. For purpose of establishing emphasis for this concept among the tape-interested members of X3T9.2, these reflections are offered.

As we all know, modern tape drives so control their data buffer and EOM early warning operations that writing to tape will cease safely ahead of the end of partition point (beyond which it is physically impossible to write). Reliable residue values can be returned by the drives in such circumstances. Typically, their buffer has been wound down and no data is left therein. If there is any, there are rules (Volume Overflow sense key (D) and possibly Recovered Buffer Data command) for dealing with the situation. This recovery concept turns out not to have been of practical importance in this situation. There won't be any data left in the buffer if operations have been normal. If an abnormality has occurred, this recovery concept is too risky to be relied upon. If writing continued to the end of partition (the old EOT) point, there is no assurance that when the medium is read--particularly if read on a different drive--that all the data supposedly written can be accessed.

The real problem to be dealt with is the case where an operating system or application software package takes advantage of the drive's ability to append data after the early warning point. If this ability is abused to the point that the drive encounters EOT, the host system should be warned emphatically that danger of data loss exists. In hindsight, the standard seems possibly to overemphasize rules for data recovery--probably a lost cause--and possibly to underemphasize the drive's responsibility to post warning to the system that it is in trouble.

Two points, then:

1. The first concern of the tape drive code writer, if EOT is ever sensed, is to immediately post a warning flag (CHECK status and sense key) to the host system. The error recovery rules are secondary; perhaps meaningless. The host will typically scrub that entire write command and start over on a new volume.

2. Since the April, 1987, plenary meeting, the sense key to flag with in this situation is SK D (Volume Overflow), not SK 3 (medium error). The basis for this decision was given in proposal 87-66. SK 3 would also provide a serious warning, but it unnecessarily confuses medium quality with improper volume termination. SK 3 primarily says to me that there is a flaw in the medium; SK D primarily says to me, "You Dummy, you screwed up!" It is useful to me as a host to have a clear distinction between these two messages. (But a number of implementations do exist which use SK 3, as you point out.)

MEMORANDUM

27 APR 1987

TO: JOHN LOHMEYER, CHAIRMAN X3T9.2

FROM: BILL SPENCE, TEXAS INSTRUMENTS

SUBJECT: TAPE WRITE EOT SENSE KEY

Based on what I have learned from you and others of the origin of Sense Key D--VOLUME OVERFLOW, based on established practice among our company and some of our vendors, and based on the simple meaning of English words, it appears that we may not have come out right in SCSI-2 when writing on tape to EOT. To correct and clarify the situation, I propose the following two corrective actions.

1. PAR 9.1.4. WRITE COMMAND (p 9-8 of draft of 10/31/86): Paragraph just before IMPLEMENTORS NOTE, next-to-last sentence. Change "MEDIUM ERROR" to "VOLUME OVERFLOW". -
2. TABLE 7-6: Sense Key Descriptions (p 7-8 of draft of 10/31/86): Sense Key Dh. Replace "remains" with "may remain".

EFFECT OF PROPOSALS

The first proposal makes the tape write response at EOT be sense key D, VOLUME OVERFLOW, instead of sense key 3, MEDIUM ERROR, in the ninth paragraph of the command, the same as it is in the seventh paragraph. (The two paragraphs cover circumstances which differ in minor degree, but to the error reporting and recovery protocols of the system this difference is of no significance.) The second proposal allows for the exceptional case where the command may just complete as the EOT indication of the medium is reached, with no buffered data.

DISCUSSION

The WRITE command case needs to be kept distinct from the READ command case. Indefinitely repeated READ commands are accepted until check status results with sense key 8 (BLANK CHECK) set, which occurs at EOD (end of data)--or, in most 1/2" reel-to-reel installations, until the system itself recognizes something in the data stream (e.g., double filemarks) which signals EOD. If before an EOD indication is returned, EOT is reached, it definitely is an error condition existing in the medium being read. In such case, sense key 3, MEDIUM ERROR, is not an inappropriate sense key, and no other sense key is close at all.

The WRITE case is quite different. In almost all implementations (every one known to us) the early warning (EOM) is sensed in plenty of time for all buffered data to be written to tape. (If this should fail, par. 7 of the WRITE command clearly--and quite properly, it seems--says that sense key D, VOLUME OVERFLOW, is to be returned. No problem.) In almost all cases, if writing continues until the drive will write no further (EOT condition), it is because the system or the user himself continues issuing write commands in the face of the warning that space is approaching an end. The error here could be a drive or firmware error; it more likely is a software or user error; but the thing that it is NOT is a medium error. In no way is it useful to report it as a medium error, which is generally recognized as an error in the magnetic storage process rather than a procedural error. But VOLUME OVERFLOW was born to deal with having reached EOT and being unable to bring a write command to an orderly conclusion. The only problem that could arise is that a completely unbuffered drive, if such still exists, would not qualify as having data remaining in its buffer--hence the proposal to recognize that in an exceptional case there may not be data in buffer despite being at EOT and unable to write further.