

Date: 07 January 1989

To: X3T9.2 SCSI-2 Working Group and Committee Members

From: Paul R. Nitza OTL Engineering
Steve Goldman Distributed Processing Technology (DPT)

Subject: Terminate Immediate with Auto Sense

Ref: X3T9.2/88-158R0 Terminate Immediate Proposal
X3T9.2/87-076R1 Auto Sense Proposal

1.0 Overview of the Problem

In some computer systems the operating system software uses main memory to cache read/write data when performing block oriented operations on I/O devices. In fact, some studies have shown that this type of caching, by the operating system, provides better performance than when the data is cached in the I/O device [Smith]. However, when caching in main memory the host cpu periodically needs to flush (especially if the cache does not perform write through) or preload the cache. This operation is easily performed under SCSI using read/write commands. However, the flush or preload operation can be quite lengthy and waiting for a normal read/write operation to complete could degrade system performance if the OS needs to do a normal I/O operation while the cache read/write is in progress. Using the Abort message would stop the read/write, but you lose any data that was transferred during the command. There is no defined mechanism in SCSI to stop an operation and determine how much of the data has been read or written before the Target stopped the operation.

2.0 Terminate Immediate with Auto Sense Proposal

This paper is a joint proposal which combines the concept of the Terminate Immediate Message (with some modifications) and Auto Sense into a single proposal to solve the above problem. The Report Transfer Length message from the original Terminate Immediate proposal, which distorted the message system, has been removed and replaced with Auto Sense. In addition a new Status byte has been added that indicates the LUN successfully processed the Terminate Immediate message.

The Terminate Immediate proposal solves the above problem by defining a new message, Terminate Immediate, that stops the current operation but sets up sense information so the Initiator can determine the number of blocks read or written before the command was terminated. Using this message the Initiator would send the message, wait for the command terminated status, wait for the command complete message/disconnect and then execute a Request Sense command to retrieve the sense information. This sequence would work but it is time consuming and another method is available that returns the sense data with the command, Auto Sense. Using Auto Sense the Initiator would send the Terminate Immediate message, wait for the command terminated status, wait for the auto sense follows message, receive the sense data and wait for the command complete sequence. Using Auto Sense completely eliminates the bus arbitration, selection and target command decode overhead of the Request Sense command.

[Smith] "Disk Cache - Miss Ratio Analysis and Design Considerations", Alan Jay Smith, University of California, Berkeley (No Date).

3.0 Terminate Immediate Proposal

This revision of the Terminate Immediate proposal consists of the Terminate Immediate message, as in the original proposal, and a new Command Terminated status. The Command Terminated status will be sent by the Target in response to a Terminate Immediate message received from an Initiator, if no other error has occurred. A new status byte is being proposed since the Check Condition status and Good status are not suitable for a command terminating in this manner. A Check Condition status indicates that an error occurred during the command, which is not the case during Terminate Immediate, while a Good status does not notify the Initiator that sense data is available.

Details of the Terminate Immediate feature and the changes needed to add it to the SCSI-2 Rev 6 document are contained in the following paragraphs.

3.1 Terminate Immediate Message

The following line must be added on Page 5-18, Table 5-2: Message Codes, after the Synchronous Transfer message:

Code	Support Init Targ	Message Name	Direction	Negate ATN Before Last Ack
11h	O O	Terminate Immediate	Out	No

The line in Table 5-2 which defines the reserved codes for single byte messages must be changed from 11h - 1Fh to 12h - 1Fh.

On page 5-32, before the section on the Wide Data Transfer Request message, the following definition for the Terminate Immediate message must be added:

5.6.xx TERMINATE IMMEDIATE Message

This message is sent from the initiator to the target to advise the target to stop the current I/O process as soon as possible, without corrupting the medium. Upon successful receipt of this message the target shall stop the current I/O process as soon as possible and terminate the command with a Command Terminated status. The sense key shall be set to NO SENSE and the additional sense code and qualifier shall be set to NO ADDITIONAL SENSE INFORMATION. If a block transfer command was being executed the sense information bytes shall be set to the first logical block address past the last logical block successfully verified, written to the medium or transferred to the initiator. If the target has completed processing the command (i.e., all data has been read, written or processed and the target was about to terminate the command) when the TERMINATE IMMEDIATE message is received it may ignore the message and terminate the command in a normal manner. If the target is unable to stop the current process it may reject this message and continue the current I/O process in a normal manner.

IMPLEMENTORS NOTE: The TERMINATE IMMEDIATE message provides a method for the initiator to force the target to reduce the transfer count of the current command to what has already been processed. The initiator can use the sense information bytes to determine the number of logical blocks that have been processed. This message will normally be used by the initiator to stop a lengthy read, write or verify operation when a higher priority command must be processed.

3.2 Command Terminated Status

The following line must be added on Page 6-9, Table 6-7: Status Byte Code Bit Values, before the Queue Full Status message:

Bits of Status Byte								Status(es) Represented
7	6	5	4	3	2	1	0	
R	R	1	0	0	0	1	R	Command Terminated

On page 6-9, the section on the Intermediate Status must be modified as shown by the italics below:

INTERMEDIATE. This status shall be returned for every command in a series of linked commands (except the last command), unless an error, exception or abnormal condition causes a *CHECK CONDITION* status, *COMMAND TERMINATED* status or a *RESERVATION CONFLICT* status to be sent:.....

On page 6-9, following the section on the Reservation Conflict Status, the following definition for the Command Terminated Status must be added:

COMMAND TERMINATED. This status shall be returned by the target after it stops the current I/O process when it has received a *TERMINATE IMMEDIATE* message (see section 5.6.xx). The *REQUEST SENSE* command should be issued by the initiator following a *COMMAND TERMINATED* status to determine the number of blocks that were processed by the target.

On page 6-17, the last sentence in the first paragraph of section 6.6, Contingent Allegiance Condition, must be changed as shown in italics below:

This guarantees that error information is available to an initiator following *CHECK CONDITION* or *COMMAND TERMINATED* status.

On page 6-17, the first sentence in the second paragraph of section 6.6, Contingent Allegiance Condition, must be changed as shown in italics below:

The contingent allegiance condition shall exist following the return of *CHECK CONDITION* or *COMMAND TERMINATED* status.

4.0 Auto Sense Proposal

This revision of the Auto Sense proposal is the same as the original one (X3T9.2/87-076R1) except it has been modified to include the new SCSI-2 changes and the Command Terminated status for the Terminate Immediate proposal.

The Auto Sense option allows the target to return sense data whenever a Check Condition or Command Terminated status is returned, without the initiator having to send a Request Sense command. If this option is enabled (a bit in the CDB control byte enables Auto Sense) the target will send an Auto Sense Data Follows message after a status phase that returned a Check Condition or Command Terminated status. If the initiator does not reject the Auto Sense Data Follows message the target will change to data in phase and send the sense data associated with the error condition. Immediately following the sense data the target will terminate the command by sending a command complete message.

Details of the Auto Sense feature and the changes needed to add it to the SCSI-2 Rev 6 document are contained in the following paragraphs.

4.1 Auto Sense Option

The following section must be added to the end of section six on page 6-23:

6.9 Auto Sense Option

The auto sense option enables the target to automatically return sense data when a CHECK CONDITION or COMMAND TERMINATED status is sent as if the initiator had sent a REQUEST SENSE command. When the target returns a CHECK CONDITION or COMMAND TERMINATED status and the auto sense enable bit in the control byte (see section 6.8) is set to one, the target shall send an AUTO SENSE DATA FOLLOWS message, change to data in phase, send the sense data and then terminate the command by sending a COMMAND COMPLETE message. Table 6-11 illustrates the phase sequence for the auto sense option. During the auto sense data phase, the target shall send all the available sense data or up to 256 bytes, whichever is less. After sending the auto sense data the target shall clear the sense data and any contingent allegiance condition as if a REQUEST SENSE command had been executed. If the initiator rejects the AUTO SENSE DATA FOLLOWS message the target shall not send the sense data and shall terminate the command in a normal manner by sending a COMMAND COMPLETE message. Under this condition the sense data can be retrieved by the initiator using a REQUEST SENSE command.

Table 6-11: Auto Sense Phase Sequence

Phase	Information Transferred to the Initiator
Status	CHECK CONDITION or COMMAND TERMINATED status
Message In	AUTO SENSE DATA FOLLOWS message
Data In	Sense Data (up to 256 bytes in length)
Message In	COMMAND COMPLETE message
Bus Free	

Byte 6, bit 7 of the Standard Inquiry Data Format, table 7-15 on page 7-19, must be changed from reserved to Auto Sense supported as shown below:

Bit	7	6	5	4	3	2	1	0
Byte								
6	ASense	Reserved						

The following paragraph must be added just before the paragraph explaining the RelAdr bit on page 7-22:

An auto sense (ASense) bit of one indicates that the target supports the auto sense option for this logical unit. A value of zero indicates the target does not support the auto sense option for this logical unit.

4.2 Auto Sense Data Follows Message

The following line must be added on Page 5-18, Table 5-2: Message Codes, after the Abort Tag message:

Code	Support Init Targ	Message Name	Direction	Negate ATN Before Last Ack
12h	O O	Auto Sense Data Follows	In	---

The line in Table 5-2 which defines the reserved codes for single byte messages must be changed from 12h - 1Fh to 13h - 1Fh.

On page 5-21, before the section on the Bus Device Reset message, the following definition for the Auto Sense Data Follows message must be added:

5.6.x AUTO SENSE DATA FOLLOWS Message

This message is sent from the target to the initiator after a CHECK CONDITION or COMMAND TERMINATED status to indicate that the target is about to return the sense data for the error condition (see section 6.xx). This message shall only be sent immediately following the status phase when a CHECK CONDITION or COMMAND TERMINATED status byte is returned and the auto sense enable bit in the command descriptor block control byte is set to one (see section 6.2.8). Upon receipt of this message the initiator shall modify the current data pointer to point to the auto sense buffer. If the target sends a RESTORE POINTERS message after the AUTO SENSE DATA FOLLOWS message the initiator shall reset the current data pointer to the start of the sense data buffer.

4.3 Auto Sense Enable Bit

One page 6-8, table 6-5: Control Byte, must be modified by changing bit 2 from reserved to Auto Sense Enable as shown in italic below:

Table 6-5: Control Byte

Bit Byte	7	6	5	4	3	2	1	0
Last	Vendor Unique		Reserved			<i>ASense</i>	Flag	Link

The following paragraph must be added on page 6-8 immediately following table 6-5: Control Byte:

ASense Bit - The auto sense enable bit (ASense), when set to one, indicates that the target shall send auto sense data immediately following a CHECK CONDITION or COMMAND TERMINATED status (see section 6.9). If this bit is set to zero, the shall not send auto sense data following a CHECK CONDITION or COMMAND TERMINATED status and the initiator must send a REQUEST SENSE command to retrieve the sense data. If the target does not support the auto sense option and this bit is set to one, the target shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB.