

X3T9.2/87-108

Date: July 10, 1986

To: John B. Lohmeyer, X3T9.2 Chair
NCR Corporation
Wichita, Kansas

From: Gary R. Stephens
IBM Corporation
Tucson, Arizona

Subject: Buffered Device Recovery after CHECK CONDITION

An unresolved item remains in Chapter 6 relative to the CHECK CONDITION for buffered devices, and in particular buffered tape devices with unwritten data blocks and/or file marks in the data buffer.

CONTINGENT ALLEGIANCE (Revisited)

At present, in SCSI-2, once a CHECK CONDITION is posted and the initiator receiving the CHECK CONDITION status requests the sense data, the target has no additional responsibility to that initiator relative to that error. The condition is also removed if a command other than INQUIRY or REQUEST SENSE is sent to the target from the initiator receiving the CHECK CONDITION. (Refer to the REQUEST SENSE command descriptions in Chapters 8 and 9.)

This architected condition of generating an implicit reservation in the device to the initiator receiving CHECK CONDITION had been defined in the sequential access device model as a contingent allegiance. This term is equivalent to that for the FIPS 65 standard covering the IBM S/370 channel. IPI-3 also has a similar situation, except that that interface provides a mechanism for an autosense function similar to that currently proposed for inclusion in SCSI-2.

The naming of the condition has been removed from the model, but the condition still remains a fundamental mechanism in SCSI-2. Note also that the SCSI-2 asynchronous event notification simulates this sequence with the first initiator to accept selection, rather than the first initiator to attempt selection.

EXTENDED CONTINGENT ALLEGIANCE (Revisited)

A problem arises in buffered tape devices after write-type operations in buffered mode when there may be several data blocks and/or file marks remaining in the data buffer which have not been written to the medium.

After the CHECK CONDITION status/REQUEST SENSE sequence or

asynchronous event notification, the target, under the present SCSI definitions, must allow commands from other initiators (other than INQUIRY and REQUEST SENSE) to be executed (assuming the initiator receiving the CHECK CONDITION does not have an outstanding RESERVE in effect). The commands from initiators unaware of the error presented to one initiator may cause a loss of position and the data buffer contents. For example, a second initiator issues the following sequence: RESERVE, REWIND, ERASE(Long), LOAD/UNLOAD(Unload), RELEASE. All position information is lost and the data in the buffer is normally purged at the rewind, thus preventing recovery by the first initiator.

If support is implemented in the target, the physical position of the error, the data blocks and the relative position and number of file marks not written to the medium may be recovered using the RECOVER BUFFERED DATA command and others. This can only occur if the first initiator is given an opportunity to execute the required commands without interference from any other initiator.

In the FIPS 65 standard, the architected condition of extending the implicit reservation to the initiator receiving the initial CHECK CONDITION while recovery is effected or until abandoned by that initiator is called extended contingent allegiance.

Again, this term was removed from the sequential access device model, but the need for defining this condition, independent of its name, is needed in SCSI-2.

PROPOSED SOLUTIONS

1. Leave the architecture as is, silent, and let various manufacturers invent mechanisms to permit recovery and have them, hopefully, document the mechanism in their publications.

This results in potentially higher attachment costs for host/initiators since the target implementation is all device dependent in nature.

2. Define that a target may present BUSY to other initiators for commands (other than INQUIRY and REQUEST SENSE) until such time that the initiator receiving the CHECK CONDITION performs the required actions to recover or executes some command interpreted by the target to abandon further recovery for the error.

Defining the blocking mechanism, but leaving the actions permitted during recovery and those actions which will signal abandoning recovery undefined still lead to considerable device dependent code in the initiator/host, but the target condition is preserved to permit recovery.

This addition seems to be the minimum acceptable addition to SCSI-2 to support the READ POSITION, RECOVER

10-51-WU,
10-51-WU,

10-51-WU,
10-51-WU,

BUFFERED DATA, and LOCATE commands added to SCSI-2 primarily for this specific condition.

3. After #2, define a flag/bit in the REQUEST SENSE data returned to the initiator which indicates that the check condition and this sense data, whether obtained via REQUEST SENSE or asynchronous event notification, have enabled the mechanism to permit the initiator to recover.

This reporting mechanism triggers the initiator/host software that is must perform some special action to recover or signal the target that it is abandoning recovery. This would also require the target manufacturers to explicitly document use of the flag/bit and the recovery mechanisms employed. Note also, that asynchronous event notification is compatible with this need since the sense data provided carries the new flag/bit.

In addition, this requires the initiator/host to perform some overt action to recover or abandon recovery. Device data integrity then is not an issue for this error should the condition of the data in the partition be discovered in error or missing at a later time.

4. After #2 and #3, add a page to MODE SENSE identifying the opcodes of the commands, which if issued, would signal the target that the initiator/host has abandoned recovery.

This last step provides a device independent mechanism to permit the initiator/host software to determine the exit conditions signaling abandoning recovery and, conversely, it also provides the menu of commands acceptable during recovery.

The initiator/host software not implemented to perform recovery now has a device independent mechanism for continuing operation with the device and freeing it up for use by other initiators.

Other mechanisms for defining which initiator SCSI IDs are attached to the same host and permitting recovery from any one of a group of initiators attached to the same host should be deferred for consideration in SCSI-3. The implementation will require either extension to the message system or new commands.

Thank you for consideration of this proposal to close out an outstanding item in the SCSI-2 draft specification.

Gary R. Stephens

226

END OF
Document,

70-51-WB1
70-51-WB1
70-51-WB1

70-51-WB1
70-51-WB1
70-51-WB1

Procl